# D5.3 Validation of Pilot Demonstrators

## Version 1.4

## Document Information

| | |
|---|---|
| Document date | 30/10/2021 |
| Contract Number | 823844 |
| Project Website | https://cheese-coe.eu/ |
| Contractual Deadline | 31/10/2021 (M35) |
| Dissemination Level | CONFIDENTIAL |
| Nature | Report |
| Coordination and editing | Tomaso Esposti Ongaro |
| Authors | Marta Pienkowska (PD1), Jorge Macías (PD2), Alice-Agnes Gabriel (PD5), Laura Sandri (PD6), Finn Løvholt (PD7), Stefano Lorito, Manuela Volpe (PD8), Vadim Monteiller (PD9), Arnau Folch (PD12); |

| | |
|---|---|
| | Joan Farnos (WP6), Josep De La Puente (WP4) and PD teams. |
| Reviewers | Finn Løvholt (NGI), Sylfest Glimsdal (NGI) |

*The ChEESE project has received funding from the European Union's Horizon 2020 research and innovation programme under the Grant Agreement No 823844*

## Change Log

| Version | Description of Change |
|---|---|
| v1.0 | Initial draft for internal working |
| v1.1 | Reviewed by the Project Coordinator |
| v1.2 | Reviewed by Pilot leaders |
| v1.3 | Reviewed by WP5 coordinator and sent out to internal reviewers |
| v1.4 | Final review by WP5 coordinator, passed to project coordinator and project manager |
| | |
| | |

# Index

# 1. Introduction

One of the main objectives of WP5 is to coordinate the effort to **make Pilot Demonstrators** (PDs, developed in WP4), **available as Services** exploitable by a broader user community. Potential users are represented by the ChEESE Industrial and User Board (IUB), but is also expected to eventually span more widely. As a measure of readiness, we demonstrate an increase in the Technology Readiness Level (TRL, defined in Deliverable D5.1) of 4 or more, for 8 of the 12 Pilot Demonstrators in ChEESE.

To achieve this goal, we have involved (in collaboration with WP6) the geophysical community (in coordination and synergy with other pan-European initiatives) and the non-academic stakeholders (industrial partners, observatories, civil protection authorities) belonging to the IUB, in the definition of the Services (i.e., what are the service objectives, which gaps have to be filled, which relevant scientific/societal questions are addressed, what are the technological bottlenecks) and in the Validation process, addressing realistic and relevant geophysical Use-Cases, with a focus on hazard and risk assessment goals.

This document reports about the main scientific and technological outcomes of the Pilot Demonstrators and the progress in the development of ChEESE Services.

- In Section 2, a summary of the methodology to **Validate** ChEESE Services and their TRL is reported.
- Section 3 summarizes the **main scientific and technological achievements** of Services and their potential impact.
- Section 4 summarizes the progress towards the **integration of ChEESE Services in the EPOS-ERIC framework**.
- Section 5 presents the **individual PD**'s scientific/technological results.
- Prototypes of Operational Services developed in this framework will be the focus of Deliverable D5.4.

# 2. From Pilot Demonstrators to Services

Pilot Demonstrators designed in WP4 aim at demonstrating that HPC codes developed and optimized in WP2, together with High Performance Workflows for data and model integration, developed and optimized in WP3, that can be exploited to address challenging scientific problems in the field of Solid Earth sciences.

WP5 aims at demonstrating advancing readiness of such applications and their added value for the scientific community and for the whole society in terms of their potential contribution to hazard assessment and risk mitigation. Eight of the twelve Pilot Demonstrators (those characterized by a target Technology Readiness Level above 4) are brought into WP5.

**Advancement of Technology Readiness Levels of ChEESE Services**

According to the definition of the ChEESE TRL scale proposed in Deliverable D5.1, we here report the advancement of TRL throughout the project development. Please note that some of the initial TRL values originally stated in the DoA have been revised in Deliverable D5.1. Most of the PDs have reached the target TRL=6, i.e. component integration and interoperability and **use-case tests demonstrated in HPC (relevant) production environments**. PD5 and PD6 engagement into PSHA and PVHA and next demonstration in an operational environment with the end-users will raise their TRL to 7-8 (next D5.4). Two Pilot Demonstrators' HPC workflows are already used in *assisted* **operational mode** (TRL between 8 and 9). In particular, the PD2 flagship code is running for faster-than-real-time tsunami simulations for ARISTOTLE service to DG-ECHO. PD12 full workflow runs daily on 1536 cores at MareNostrum-4 Tier-0 machine at BSC to advise regional airport management and air traffic (e.g. during the volcanic eruption at La Palma island in the Canary archipelago).

| Pilot | PD1 | PD2 | PD5 | PD6 | PD7 | PD8 | PD9 | PD12 |
|---|---|---|---|---|---|---|---|---|
| **Initial TRL** | 3 | 3 | 4 | 3 | 3 | 3 | 4 | 3 |
| **Target TRL** | 5-6 | 6-7 | 6 | 7 | 5-7 | 6-8 | 6 | 6 |
| **Achieved TRL** | **5** | **8-9** | **6 (7-8)** | **6 (7-8)** | **6 (7-8)** | **7-8** | **6** | **8-9** |

*Table 1. Summary of Pilot Demonstrators' TRL achieved with Task 5.3. In parentheses, the expected TRL to be achieved within Task 5.4.*

**Criteria for quality assessment and validation**

**Quality assessment** of the delivered products is ensured by the criteria of *peer-review* and scientific publication. Every Pilot Demonstrator reports the list of publications and dissemination in the scientific community. The total number of XXX publications on peer-reviewed journals and YYY conference presentations related to Task 5.3 and reported in this Deliverable demonstrates such a commitment.

**Validation**, on the other hand, is a *context-specific* assessment about whether previous *gaps* have been filled by the PD development. In this Deliverable, this has been done through a *self-assessment*, based on the achievement of previously defined target *functional requirements*. Such requirements had been identified (on the basis of the end users needs) and reported in Deliverable D5.1. In Deliverable D5.4, an *external assessment* will be done through the direct involvement of the end-users in Pilot Demonstrator Exercises.

**Data and model interoperability**

Pilot Demonstrators made a step forward towards standardization of the digital formats of both output data and associated metadata.

For data input and output, most of the PDs adopted a standard, machine-independent, self-describing Hierarchical Data Format (HDF5) supporting multidimensional arrays of scientific data. The data format has mostly been implemented by using the NetCDF libraries (https://www.unidata.ucar.edu/software/netcdf/). These data formats allow efficient performance when including I/O, up to sustained petascale (e.g., Krenz et al., SC21), in PD production runs.

To ensure model interoperability, a collaborative effort among PD5, PD6, PD7, PD8 and PD12 has been aimed at defining two new standard formats for Probabilistic Hazard data. Such formats did not exist before and they will be shared with the geophysical community and in particular with the EPOS TCSs.

Both formats follow the general CF (https://cfconventions.org/) conventions, and all variables and dimensions should follow a specific naming convention. The formats differ on the choice of the reference quantity used as dimension of the NetCDF file.

The preferred format is "Format Hazard-maps" and it uses as reference value the probability values. This format allows the user to store, for each probability value, different hazard intensities. If epistemic uncertainty is computed, the intensity corresponding to different percentiles can be stored as well. This format is suitable for storing multiple hazards in the same file, since probability values are in common to all hazards.

The secondary format is "Format Probability-maps" and it uses as reference value the hazard intensity values (e.g. Peak Ground Acceleration, Ground Ash Load, etc.). This format allows the user to store, for each intensity value, the corresponding probability value. If epistemic uncertainty is computed, the probability corresponding to different percentiles can be stored as well. This format is more suitable for single-hazards, since hazard intensity is dependent on the specific problem under investigation.

For what concerns Metadata, their semantics have been discussed among PDs teams and some of the stakeholders (e.g., IMO) and will be submitted to peer-review on internationally recognized journals and scientific conferences. Their format and structure is designed to be compatible with EPOS standards.

**Involvement and role of the end-users**

In the transition from TRL<=4 (laboratory) to TRL>=5 (production), end-users (including members of scientific associations and recognizable scientific communities) have been involved in the co-design of Services by defining, in synergy with the ChEESE teams and the PD leaders, the **appropriate use-cases** and **requirements for the validation** (Deliverable D5.1).

For Deliverable D5.3, most of the effort has been dedicated to the involvement of the **scientific community** in the discussion of the scientific context and sharing of the computational tools and results. Some of the **IUB members** and other stakeholders have been deeply involved in the development of use-cases and deployment of the demonstrators, in particular: Icelandic Meteorological Office (PD5, PD12), Spanish, Norwegian and Argentinean Meteorological Services (PD12), the Italian Civil Protection Department (PD6, PD7, PD12), ARISTOTLE (PD2, PD6), the SiAM Italian National Tsunami Warning Centre (PD7, PD8), Air traffic regulators (ISAVIA in Iceland and ENAIRE in Spain - PD12). These are detailed in each individual PD report (Section 5).

Deliverable D5.4 will have a more direct involvement of end-users for the development of operational prototypes.

## 3. Summary of achievements

ChEESE services are based on specific Workflows (see WP3) customized to solve the use-cases and optimized to run with the ChEESE Flagship codes (see WP2) **SALVUS** (PD1,PD9), **Tsunami-HySEA** (PD2, PD7, PD8), **SeisSol** (PD5), **Fall3D** (PD6, PD12), **SPECFEM3D** (PD5, PD12).
Each PD report in Section 5 summarizes the scientific and technological advancements obtained in the integration of the optimized numerical engines into complex workflows, including pre- and post-processing of numerical results to achieve the scientific goals.

**Exploitation of HPC architectures**
The exploitation of HPC supercomputers differs for the different applications:

- **HPC Capability**: running a single, or a relatively small number of tightly coupled simulations on a high number of computer nodes to achieve the desired resolution/extension/size of the physical problem
- **HPC Capacity**: Running ensembles of many ~petascale simulations, that aggregate to exascale by exploiting complex workflows, including interoperability of models and data, on a single exascale machine.

ChEESE Urgent Computing workflow (PD1) mostly exploits HPC Capability, whereas Probabilistic Hazard Assessment (PD6, PD7) workflows mostly exploit HPC Capacity. Probabilistic Forecasts and Early Warning (PD5, PD8, PD12) and Seismic Tomography (PD9) all exploit a combination of Capacity and Capability. Finally, PD2, based on optimized GPU Urgent Computing, is also one of the building bricks of PD7 and PD8.
WP5 use-cases have exploited the computing capability of the following machines, mostly through PRACE grants:

- (PD7,PD8,PD9) Marconi100, the new CINECA accelerated Tier-0 cluster based on IBM Power9 architecture and Volta NVIDIA GPUs (https://www.hpc.cineca.it/hardware/marconi100).
- (PD1, PD5, PD12) Mare Nostrum 4, the BSC Tier0 cluster based on Intel Xeon chip (https://www.bsc.es/marenostrum/marenostrum)
- (PD6, PD12) Joliot-Curie, the TGCC BULL Sequana X1000 Petascale System.(http://www-hpc.cea.fr/en/complexe/tgcc-JoliotCurie.htm)
- (PD5) SuperMUC-NG, the LRZ Petascale System (https://doku.lrz.de/display/PUBLIC/SuperMUC-NG)

## 4. Integration within the EPOS-ERIC

ChEESE has interacted with EPOS (European Plate Observing System) ERIC (European Research Infrastructure Consortium) towards the common aim of providing the geophysical community with **Software and Workflows as Services (Saas and WaaS)**. While this requires an engagement by EPOS that is beyond the scope of current activities, a common roadmap is being developed towards fulfilling this aim, while several activities are being already performed thanks to the interaction with the EPOS

Sustainability Phase Project (EPOS SP), or with EPOS activities funded at the national level.

The first steps in this direction have included:

- Development of a new EPOS Tsunami TCS (Thematic Core Service), namely the TCS Tsunami, which has recently achieved the formal status of candidate. The peculiarity of the TCS Tsunami is a strong computational-science oriented character stemming from the traditionally large use of numerical simulation by the tsunami community; some tsunami-related ChEESE PDs (PD2, PD7, PD8 will soon be accessible through the portal of the Italian national node currently under construction (https://www.tsunamidata.org/).
- In preparation of the future distribution of several SaaS and WaaS which will involve the products from several ChEESE PDs (PD2, PD4, PD5, PD6, PD7, PD8, PD12), several actions are being performed in close interaction with EPOS, such as
    - Addressing the compatibility of the data and metadata provisional standards being defined by ChEESE with the EPOS ones.
    - Participation in the co-design of the future EPOS ICS-D (Distributed Integrated Core Services).
- Provision of HPC use cases, in particular from the Geohazard Permanent Supersites (https://geo-gsnl.org/).
- Collaborations for training of European scientists and dissemination of ChEESE results through EPOS.

# 5. PD sheets

The next sections describe the development of HPC Services for 8 Pilot Demonstrators, targeting TRL 5-8. In particular, it reports the results of the target Use Cases, and the Validation in terms of Functional/non-Functional Requirements met. This description will be complemented by information collected in the ChEESE Pilot Service prototype and enabling (Deliverable 5.4, 31 January 2022).

## PD1. Urgent Seismic Computing

| PD1 | Urgent Seismic Computing |
|---|---|
| Leader | Marta Pienkowska (ETH), Andreas Fichtner (ETH) |
| Participants | Marisol Monterrubio-Velasco (BSC), Juan E Rodríguez (BSC), Otilio Rojas (BSC), Josep de la Puente (BSC) |
| Workflow | UCIS4EQ |
| Engine | SALVUS |
| TRL initial | 3 |
| TRL target | 5-6 |
| TRL achieved | 5 |

**HPC Products (available software and workflows)**

Within PD1 of ChEESE we have developed a prototype of an HPC-based urgent seismic simulation workflow - the **Urgent Computing Integrated Services for EarthQuakes** (UCIS4EQ) - delivered as a collection of virtual services. UCIS4EQ aims to **rapidly deliver information on the distribution of shaking intensities of moderate to large earthquakes** and complement the existing early impact assessments with high-resolution synthetic data. To this end, the ChEESE flagship code Salvus has been integrated within UCIS4EQ to perform the HPC simulations based on the workflow-generated input parameters.

Within UCIS4EQ we distinguish between two systems - a future front-end and a back-end that consists of a set of workflow-manager-orchestrated processes. The front-end user interface remains to be implemented, while the back-end services include:

(1) Event Detection: an automatic service external to the UCIS4EQ workflow that continuously queries external web services (provided by FDSN, the International Federation of Seismograph Networks) in order to identify candidate events for an urgent simulation.
(2) Urgent Computing workflow: a set of processes where micro-services act as access points for the execution of each given task. The workflow manager coordinates the inputs and outputs of each task, including submitting and monitoring the job submission on HPC clusters. Multiple instances of the service can be run simultaneously, in case that more than one event fulfills the requirements for an urgent simulation. It should be noted that some components of UCIS4EQ use third-party software that has been integrated into the workflow, such as the Graves-Pitarka rupture generator or the Salvus wave propagation software suite.

(3) Data Access Layer (DAL): an interface for interacting both with short-term (intermediate) information and with long-term data. The short-term information is generated both by the Event Detection and the Urgent Computing workflow components and is stored in non-relational databases (such as event data or the state of the execution of the workflow or execution metadata). The long-term data (such as seismic velocity models, simulation meshes, off-line trained models, historical past earthquake events information) is stored on services such as EUDAT B2SAFE and the DAL enables read and write access through appropriate communication and authentication protocols.

The workflow allows for integrating a range of wave propagation software and communicates with the software via YAML files that specify the required simulation parameters. SALVUS is the ChEESE flagship code that is now the PD1 production software that has been fully integrated into UCIS4EQ and tested, while the other codes remain to be tested in an operational environment.

**Use case #1. Mediterranean: Samos-Izmir 20.10.2020 earthquake**

***Summary of technological achievements (HPC performance, etc)***

In this use case, developed towards task T3.5 Pre-exascale testbed execution, **we have run the largest SALVUS simulation to date** which was accurate up to 20 Hz (dominant frequency of 10 Hz) and which required **the use of the whole BSC-CNS Marenostrum 4 supercomputer.** We have therefore challenged the main technological bottlenecks related to the scaling-up of the problem, and we have identified the key improvements that are necessary for the functioning of the service at scale:

- Mesh generation. We have tested the automatic generation of unstructured meshes that include both topography and bathymetry for over 215 million elements, with 512 grid points per element. Generating such meshes required 250 GB of RAM, which will be a significant bottleneck for routine applications at large scales. Chunk-wise mesh generation - saving parts of a generated mesh of a limited size to disk - is a solution currently under development.
- Mesh reading and decomposition. The new parallel I/O strategy that enables the parallel reading of large unstructured meshes (implemented within ChEESE, see WP2) has been tested at an unprecedented scale. Without such strategy, most of the runs scaling up to the 20 Hz resolution would not have been achievable. A careful consideration and tuning of available memory per process was necessary both for reading and decomposing the meshes.
- Finite source representation: With the increase in frequency, the kinematic representation of the fault results in a large number of point sources. The current implementation in Salvus has proven to not be optimal for such use cases and will require further work. Moreover, it highlighted issues with load balancing that were not evident for smaller scale scenarios.
- Output management: Densely sampled output has proven challenging due to required RAM memory, dump times and load balancing issues. This prompted

considerations of on-the-fly processing of outputs to mitigate this for future applications and maintain the desired scaling. A fast Fourier transform (FFT) can already be performed on-the-fly to generate frequency domain outputs for specific discrete frequencies, and similar solutions should be sought for ground motion proxies.

- Checkpointing: Currently not available in Salvus, and a feature required for UCIS4EQ to mitigate node-failures and other hardware issues.

***Scientific achievements***

Working towards the scaling up of realistically set-up simulations we have shown the feasibility of large-scale simulations that cover the frequency range of interest. Pure performance benchmarks involve simplistic set-ups (e.g. regular cube meshes and a single point source) that often do not represent the requirements of production-grade scenarios (e.g. including heterogeneous models and geometric features such as topography and bathymetry, with complex source representations and densely sampled outputs) and thus do not expose practical bottlenecks that impact the outputs of the service.

In PD1, the key functional requirements of a future urgent computing service are the time to solution and the frequency resolution of the synthetic ground motions, which are interdependent. The simulated maximum frequency is related to the time to solution and the computational resources available, as longer time to solution and more resources allow for higher frequency content. It is therefore a balance to strike for each specific use case, as ground motion intensities are strongly dependent on frequency. For UCIS4EQ it is key to explore and test the entire frequency range of interest for structural engineering (up to 10 Hz). It was a major milestone to work towards realistic scenarios that cover the upper bound of the scales of interest and thus to prove the feasibility of the urgent computing service.

***Scientific Products***

***Conferences, seminars***

- de la Puente, Josep; Rodriguez, Juan Esteban; Monterrubio-Velasco, Marisol; Rojas, Otilio; Folch, Arnau; Urgent Supercomputing of Earthquakes: Use Case for Civil Protection Proceedings of the Platform for Advanced Scientific Computing Conference 1-8 2020
- Monterrubio-Velasco, Marisol; Carrasco-Jimenez, José Carlos; Rojas, Otilio; Esteban Rodríguez, Juan; de la Puente, Josep; Fast acquisition of focal mechanism based on statistical analysis EGU General Assembly Conference Abstracts 3224 2020
- Folch, Arnau; de la Puente, Josep; Sandri, Laura; Halldorsson, Benedikt; Pienkowska, Marta; Gracia, Jose; Lanucara, Piero; Bader, Michael; Gabriel, Alice-Agnes; Macias, Jorge; Preparing Earth Sciences to Upcoming Infrastructures. The Center of Excellence for Exascale in Solid Earth (ChEESE) AGU Fall Meeting Abstracts 2020 IN003-01 2020

- Pieńkowska, Marta; Esteban Rodríguez, Juan; de la Puente, Josep; Fichtner, Andreas; Deterministic modelling of seismic waves in the Urgent Computing context: progress towards a short-term assessment of seismic hazard EGU General Assembly Conference Abstracts EGU21-15516 2021
- Monterrubio-Velasco, Marisol; Carrasco-Jimenez, J Carlos; Rojas, Otilio; Rodriguez, Juan E; Modesto, David; de la Puente, Josep; Source Parameter Sensitivity of Earthquake Simulations assisted by Machine Learning EGU General Assembly Conference Abstracts EGU21-5995 2021
- Pienkowska, Marta; Urgent Computing for Seismic Hazard Assessment: Progress and Challenges; part of the ChEESE CoE European Urgent Computing Workshop on the EuroHPC Summit Week 2021

**Validation.**

| **Functional requirement** | **Target (from D5.1)** | **Achieved** | **Validated (YES/NO)** |
|---|---|---|---|
| Time to solution. | A few hours to a day. | 1h for 1Hz dominant period.<br>4h for 2.5Hz dominant period.<br>12h for 10Hz dominant period. | YES |
| Resolution. | Targeting above 1Hz with a maximum of about 10Hz. | 10 Hz dominant period. | YES |
| Number of simulations. | From single high frequency runs to small ensemble runs. | Single high frequency runs (as above), small ensemble runs of lower frequency (2.5Hz dominant period). | YES |
| Data formats. | No exact data formats were targeted. | Raw data in HDF5. Processed data as images and in HDF5. | YES |
| Uncertainty quantification. | No specific uncertainty treatments were targeted. | We are still working towards performing ensemble runs for UQ. | n.a. |

Table 1.1. Validation criteria for PD1.

**Involvement of end-users**

We are working with Global Parametrics (GP) to further refine the use case for the Mediterranean region and explore simulation-based parametric financing. Currently GP uses basic past and stochastic source information (location and magnitude) for simple pricing and payout models globally. We are working towards simulating both the past and the stochastic events for the Mediterranean region to generate 1Hz synthetic datasets and explore pricing and payout models beyond just the source parameters, but accounting for spatial variability based on simulated ground motion intensities.

We are also at the stage of preliminary discussions with the Swiss Seismological Service (SED) - the federal agency responsible for monitoring earthquakes - regarding a systematic comparison and cross-validation of synthetic results with the current state-of-the-art shakemaps. This is the first step towards synthetic ground motion intensities becoming a complementary product to the available GMPE- and data-based shakemaps.

**Impact**

In PD1 we have proposed a prototype of a seismic urgent computing workflow that automatically generates synthetic ground shaking information following an event trigger. Given basic information from the alert system, the workflow estimates source parameters and defines other simulation inputs, and finally it launches a deterministic simulation of the propagation of seismic waves. At TRL level 5 the service remains targeted at field specialists and needs to further be tested in a controlled environment for specific use cases, but it demonstrates the feasibility of urgent computing for seismic applications. It therefore opens an avenue towards rapid synthetic ground motion maps that could complement the current approaches. Moreover, UCIS4EQ is now also part of an effort to define the required HPC urgent computing protocols that would open HPC infrastructures to urgent computing (in collaboration with the Poznan Supercomputing and Networking Center (PSNC) via the EuroHPC funded eFlows4HPC project).

## PD2. Faster Than Real Time Tsunami Simulations (FTRT)

| PD2 | Faster Than Real Time Tsunami Simulations (FTRT) |
|---|---|
| Leader | Jorge Macías (UMA) |
| Participants | M.J. Castro, M. de la Asunción, C. Sánchez Linares, J.M. González-Vida, INGV, NGI |
| Workflow | *TS-Workflow;* |
| Engine | *Tsunami-HySEA* |
| TRL initial | 3 |
| TRL target | 6-7 |
| TRL achieved | 8-9 |

**HPC Products (available software and workflows)**

The ChEESE PD2, Faster Than Real Time (FTRT) Tsunami Simulations, consists in a workflow for performing single, multiple or massive numerical simulations, depending on the computational resources available but always faster than real time, meaning this at shorter times than the natural evolution of the real tsunami itself (T4.2.1). This PD also aims to advance in Urgent Computing (T4.2.2) and Capacity Simulations (T.4.2.3) as subproducts of the same workflow, the last one being key for the two related tsunami PDs, PD7 (PTHA, Probabilistic Tsunami Hazard Assessment) and PD8 (PTF, Probabilistic Tsunami Forecast). In deliverable D4.2 a description of PD2 can be found as well as several testing examples.

ChEESE PD2 relies on the *Tsunami-HySEA* code. *Tsunami-HySEA* is a numerical model to simulate tsunamis generated by earthquakes, where the seafloor deformations caused by earthquakes are obtained using the Okada model. Two different versions of *Tsunami-HySEA* have been developed during the ChEESE project: a multi-GPU version that performs one single simulation distributed on multiple GPUs, and a Monte Carlo version (coined *Tsunami-HySEA MC*) that performs multiple simulations in one single execution, where each simulation is performed on a single GPU. All of them include nested meshes processing, activation of the nested meshes processing, asynchronous file writing, and time series saving.

Here we summarize some important versions of the *Tsunami-HySEA* multi-GPU code that have been released during the ChEESE project:

- *Tsunami-HySEA 3.4.0*: This was one of the earlier versions of Tsunami-HySEA in the ChEESE project. It has the features cited in the previous paragraph.
- *Tsunami-HySEA 3.6.1*: This version adds variable friction, a sponge layer to avoid numerical instabilities in some cases, a new algorithm to process the

nested meshes for inundation simulations, and an improved efficiency as a result of the first code audit.

- *Tsunami-HySEA 3.7.1*: This version adds the possibility of reading Okada earthquake source information from an external file, different saving times for each resolution level, the storage of the NetCDF data in single precision except the longitude and latitude, and it outputs warning messages if a particular Okada deformation has unexpected values.
- *Tsunami-HySEA 3.8.1*: This version adds two more metrics that can be stored (maximum modulus of velocity and maximum mass flow), and the use of the int64_t datatype for the size of many arrays so that bigger meshes can be used.

Regarding the *Tsunami-HySEA MC* version, next we summarize some important versions of this code:

- *Tsunami-HySEA 3.8.1 MC*: This version has the same general features as the multi-GPU 3.8.1 version. For the time series it stores the water height and the velocities at each time step for all the points. All the saved data is compressed using the deflate feature of NetCDF.
- *Tsunami-HySEA 3.8.4 MC*: This version removes the storage of the velocities in the time series, and packs the water height in a short int to reduce its size to the half in the NetCDF files. It also slightly improves the efficiency as a result of the second code audit, by changing some instructions in one CUDA kernel.
- *Tsunami-HySEA 3.8.5 MC*: This version adds the possibility of ending the execution if a certain percentage of simulations fail. It also improves the efficiency by commenting out all the parts of the code corresponding to the message passings that are not performed in the MC version.

*Scripts and tools*

- (for use-case #1). Development of a message passing system based on RabbitMQ to be used between the SPADA system and the supercomputing services . Depending on the earthquake epicenter location the system is able to automatically select an optimal computational grid size and refinement level depending on the seismological parameters. Preprocessing tools: Automatic grid generator, parameter file and launch scripts. Postprocessing tools: Automatic graphical output scripts according to the visualization scale specifically designed for the TWS. System operation control scripts: Scripts to control side by side if the system is running properly.
- (for use-case #2). Scripts and tools to generate the topobathymetric files and initial conditions. Tools for generating nested grids at a given resolution. Preprocessing tools (bash-scripts) to efficiently group the simulation runs according to the resources available in the HPC. Postprocessing tools for visualization and to georeference the output files that allow the elaboration of maps where the variables resulting from the simulation are reflected. These tools have been optimized and migrated to python scripts for better integration in HPC environments.

### Use case #1. Tsunami Service for Aristotle (Urgent computing)

#### *The ARISTOTLE project*

The urgent computing (UC) capabilities in the tsunami natural hazard framework are strengthening the monitoring and analysis functions of the European Emergency Response CoordinationCentre (ERCC) and its Situational Awareness Sector (SAS) by helping to design the multi-hazard advice service at global level and on a 24/7 operational basis. In this context, the ARISTOTLE-eENHSP project (All Risk Inte-grated System TOwards Trans-boundary hoListic Early-warning - enhanced European Natural Hazards ScientificPartnership) has been designed to offer a flexible and scalable system that can provide new hazard-related services to the ERCC. The ARISTOTLE consortium includes 18 partner institutions operating in the Meteorological and Geophysical domains. It builds on a consolidated and multi-disciplinary partnership consisting of world-leading scientific centres in the areas of Earth and Climate sciences, providing operational and monitoring services, early warning and information systems as well as contributing to innovation and research actions. The ARISTOTLE-eENHSP Consortium is currently providing advice in a multi-hazard fashion for the following inter-related hazards: earthquakes, tsunamis, volcanic eruptions, severe weather events, flooding events and wildfires/forest fires.

#### *Summary of technological achievements (HPC performance, etc)*

The EDANYA Group of the University of Málaga has participated in this consortium since October 2019 providing services of Urgent Computing through Faster than Real Time (FTRT) tsunami computations with the numerical model *Tsunami-HySEA*.

The ARISTOTLE tsunami service (TS) is integrated in the SPADA (Scientific Products Archiving and Document Assembly) IT platform that gathers the scientific, exposure and preliminary impact information which are used by the multi-hazard operational board to assemble the reports. This platform relies on existing and newly developed web services. The **TS workflow** (Fig. 2.1) consists of four steps: the system is triggered by an end-user who is on duty in the service. Using the earthquake parameters that can be provided by different seismic monitoring sources, the scenario parameterization is defined and it is sent to the supercomputation resources (in this case located at INGV and the University of Málaga). In this step a message passing system (RabbitMQ) is used between the SPADA system and the supercomputing services. Depending on the earthquake epicenter location the system is able to automatically select an optimal computational grid size and refinement level depending on the seismological parameters. For example, as it will be presented later, if the epicenter is located in the Mediterranean Sea, the system automatically performs 8 hours of wall clock simulation in a 2 arc-min resolution grid, then detects the limits where the tsunami has arrived and later performs a second simulation in a new domain with more resolution (30 arc-sec). Depending on the event, the computation time can last from a few seconds to the order of minutes. The current outputs of the tsunami service system are: maximum water elevation in the considered domain, wave arrival time and maximum water elevation along coast locations.

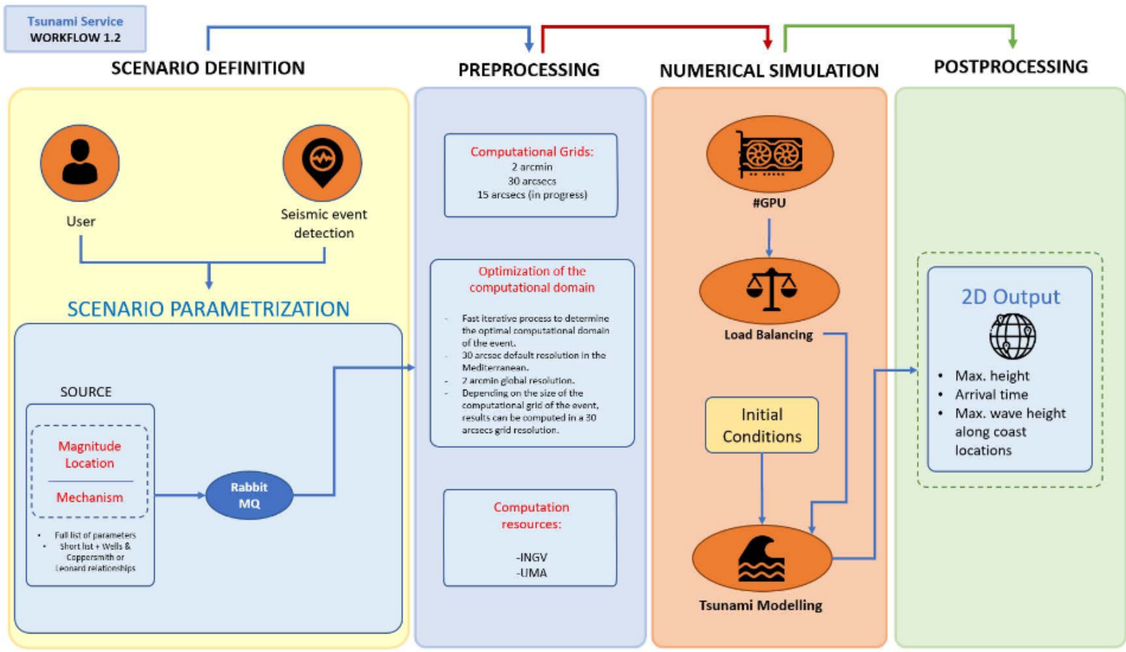


Figure 2.1. Diagram representing the tsunami service (TS) Workflow

This procedure is scalable depending on different aspects, like the computation resources or the Digital Terrain Models (DTM) available. Consequently, the numerical computation output could be improved in different ways: for instance by improving the grid resolution (even using nested meshes in specific areas of interest), or even providing not only one scenario output but considering an ensemble of cases that serves as input to the the Probabilistic Tsunami Forecast (PTF), PD8 (see below).

The system outputs are delivered to the European Emergency Response Coordination Centre (ERCC) in a multi-hazard report providing expert analysis made by an expert panel in the different involved hazards. In our case, the sunami service outputs are relevant in the sense that they have to be easily comprehended by end-users. For instance, an enhanced semaphore colorbar has been designed where each semaphore color: green, yellow, orange and red has been subdivided into three sub-colors. The output is clear for end-users providing some basic information (Fig. 2.2, left panel). The arrival times figure output has also been improved with the addition of a jet colormap that completes the information given by the isochrons.

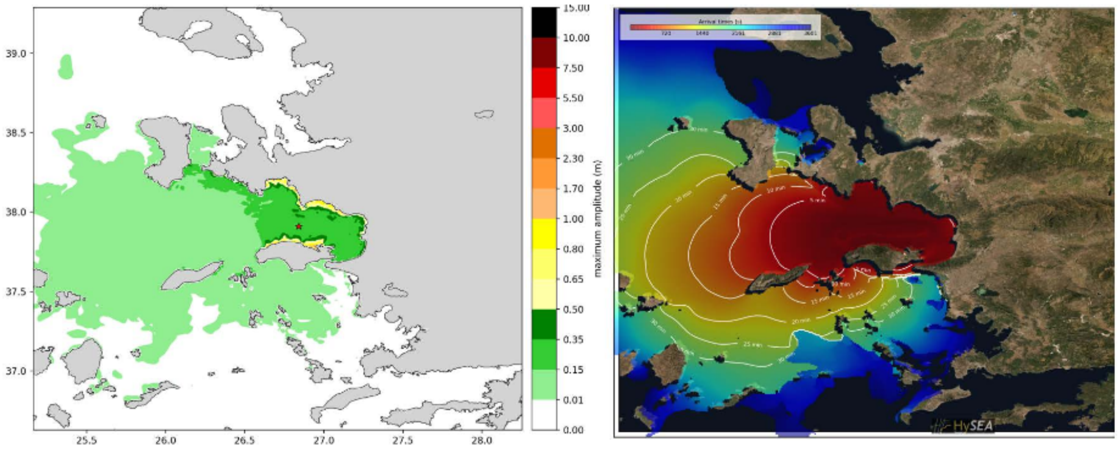


Figure 2.2. Graphic representation of numerical simulation outputs from the TS workflow.

To illustrate the computational efficiency of the system, in Tab. 2.1 are shown the computation times required to perform the process described in the previous section. In this case it is simulated one hour of wall clock propagation of the tsunami in the 2 arc-min resolution and then the same time in a 30 arc-min resolution grid adapted to the domain where the tsunami waves have arrived. The total time is 38.48 seconds in a single P100 NVidia GPU.

The EDANYA group role in this project is related to the Tsunami Service with the development and tuning of the system. ***In fact, the 8-9 technology readiness level (TRL) achieved with Tsunami-HySEA in this service makes it operational*** and at the same time scalable to incorporate the new state of the art techniques when they are available.

| **Scenario** | **Resolution** | **Wall Clock simul. time** | **N. cells** | **Comp. Time** | **Postproc. time** | **Total time** |
|---|---|---|---|---|---|---|
| Izmir. | 2 arc-min | 60 mins (3600 secs) | 1305x480 = 626400 | 1.15 secs | | |
| Izmir | 30 arc-sec | 60 mins (3600 secs) | 460x336=154560 | 2.22 secs | 35.11 secs | 38.48 secs |

Table 2.1. Computational performance for use-case #1.

***Scientific achievements***

The present use case is an operational service itself and it is functioning using the code *Tsunami-HySEA* since the 1st October 2019. It represents a major step forward in tsunami emergency management for the ERCC, which can foresee the potential humanitarian aid that will eventually be required from affected countries anywhere on the planet.

***Scientific Products***

***Conferences, seminars***

- J.M. González-Vida, S. Ortega, J. Macías, M.J. Castro, A. Michelini and A. Azzarone. What is the humanitarian aid required after a tsunami?. XXVI CEDYA/ XVI Congreso de Matemática Aplicada, Gijón (Spain), 14-18 June 2021. XXVI CEDYA/ XVI CMA Proceedings. ISBN 978-84-18482-21-2, pp. 197-200, http://hdl.handle.net/10651/59084

**Use case #2. Multi-scenario simulations for the Spanish TEWS (Early warning)**

This use case has been directly proposed by one of the members of the ChEESE project User Board, the IGN (Instituto Geográfico Nacional, Spain). The specific problem proposed and configuration the IGN is asking us to implement as a use case is of their interest to be applied and used in the National TEWS in Spain. Here, in particular, we deal with a tsunami event in the Gulf of Cádiz with main impact in Huelva and Cádiz (Andalusian Atlantic Coast) but also in the Canary Islands and Galice (North-Western Spain).

The idea beneath the use-case proposed is very simple, it consists in considering some variability in the seismic source as defined by means of the Okada model parameters. More precisely, once the "reference" event has been defined, instead of simply using those parameters to perform a single deterministic simulation, some of them are varied and multiple simulations are performed. Thus, after seismic detection a location and magnitude are obtained (lat, lon, depth, Mw), with this information, a fault from a Database of known active fault is automatically selected and the associated Okada parameters extracted (every location is associated to a particular fault). This is the "reference" event or parameters. Then uncertainty in the source definition is considered varying location (moving N, E, S, and W) the epicenter, magnitude, strike and dip, up to consider a total of 135 close scenarios to the "reference" event.

In order to give a faster and progressively more accurate warning response, simulations will be carried out at different computational domains and resolutions in clusters of consecutive 135 simulations, in a certain order, in such a way that the alert messages can be updated as the simulations are completed at an increased resolution.

The output variables to analyse are maximum wave amplitude of the time series at the FCP (Focal Points) along the Spanish coast for each simulation. These FCPs are defined and provided by the IGN and are used as reference wave height locations in the Spanish Warning System. Arrival times and maximum wave amplitude in the whole domain is also provided.

The final output are alert level maps corresponding to the mean of the maximum height at each FCP clustered by coastal segments. The numerical results obtained can be also used a posteriori to assess the sensibility of the results to the varying parameters. In a few minutes an alert level along the different segments along the Spanish coast (in the present use case in Huelva y Cádiz) will be provided and certain variability/uncertainty will be included.

### *Scientific achievements*

The use case presented here is aimed to have a direct applicability in the Spanish National Tsunami Warning System, in the operational TEWS implemented by IGN. It represents a qualitative step forward in current systems in the NEAM region as they are based in decision matrices and precomputed databases, but do not use real time on the fly computations, in the present use case including some uncertainty in the source definition.

### *Scientific Products*

#### *Published papers.*

- F. Løvholt, S. Lorito, J. Macías, M. Volpe, J. Selva and S. Gibbons, "Urgent Tsunami Computing," 2019 IEEE/ACM HPC for Urgent Decision Making (UrgentHPC), Denver, CO, USA, 2019, pp. 45-50, doi: 10.1109/UrgentHPC49580.2019.00011.
- J.M. González-Vida, M., Castro, J. Macías, M. de la Asunción, S. Ortega, and C. Parés (2021) Tsunami-HySEA: A Numerical Model Developed for Tsunami Early Warning Systems (TEWS). In: Cruz M., Parés C., Quintela P. (eds) Progress in Industrial Mathematics: Success Stories. SEMA SIMAI Springer Series, vol 5. Springer, Cham. [doi: 10.1007/978-3-030-61844-5_12].

#### *Conferences, seminars*

- J. Macías and M. de la Asunción (2019). “Where are the limits? (In FTRT tsunami computations).” EGU2019-11357 NH5.1/OS2.22/SM3.11. European Geoscience Union General Assembly (EGU 2019). Geophysical Research Abstracts, Vol. 21, EGU2019-11357, 2019.
- J. Macías et al. “Faster Than Real Time tsunami simulations: challenges and solutions towards High Performance Exascale Computing”. European Geoscience Union General Assembly (EGU 2020). Geophysical Research Abstracts, EGU2020-19848, 2020. Session: NH5.1 – Tsunamis. Viena (Austria), virtual conference, 3-8 May 2020. url: https://meetingorganizer.copernicus.org/EGU2020/EGU2020-19848.pdf doi: 10.5194/egusphere-egu2020-19848
- S. Gibbons, F. Lovholt, S. Lorito, J. Macías et al (2021). “HPC for urgent tsunami computation”. Keynote talk in virtual EuroHPC Summit Week 2021. Session: European Urgent Computing workshop. url: https://events.prace-ri.eu/event/1018, 22-26 March 2021. 24 March.
- C. Escalante, M.J. Castro and J. Macías. “Dispersive Tsunami-HySEA model for Faster Than Real Time tsunami simulations”. The PASC21 Conference (Platform for Advanced Scientific Computing). Advances in Computational

Geosciences, Part III. url: https://pasc21.e-event.ch/minisymposium/msa343. University of Geneva, Switzerland / Virtual, 5-8 July 2021. 6th July.

**Validation.**

| **Functional requirement (these are examples)** | **Target (from D5.1)** | **Achieved** | **Validated (YES/NO)** |
|---|---|---|---|
| Time to solution. | Less than 10 minutes | We achieve the first alert level on a coarser global mesh in less than 1 minute then, we continue updating the alert levels using three reduced in size but increasing the numerical resolution and this in less than 7 minutes | YES |
| Domain size and resolution. | Go from 1/2 arcmin to 1/16 arcmin in order to give the alert level with greater precision as we solve simulations with better resolution. | Grids from 4M up to 9M volumes depending on the resolucion for each simulation. | YES |
| Number of simulations. | 135 simultaneous simulations | 135 simultaneous simulations | YES |
| Data formats. | NetCDF output files | NetCDF | YES |
| Uncertainty quantification. | Inclusion of uncertainty in 4 of the parameters that determine the initial conditions of the event. | Statistical measures of the results | YES |

Table 2.2. Validation criteria for PD2.

**Involvement of end-users**

As already mentioned, this particular use case has been proposed by IGN and the actual set-up has been co-designed by the IGN and ChEESE team. For defining the final set-up we have provided computing times for different computational domains and using different mesh resolutions. Finally, four domains covering different areas at 3 resolutions have been retained and computations performed at different levels of priority. Several virtual and in-person meetings have taken place between IGN and UMA researchers.

**Impact**

In 2018 we were awarded with the NVIDIA Global Impact Award for our contribution to improve TEWS using GPUs. Now we are going much further really using HPC capacities to boost tsunami simulations in Early Warning, Urgent Computing, PTHA and PTF, and hence as input to the downstream PDs 7-8. In particular, the present use case represents the more modest of these achievements but a qualitative step forward in state-of-the-art TEWS in the NEAM (North Eastern Atlantic and Mediterranean Sea) region. We are not only including computational solutions but also considering some variability in the provided assessment. This is something that is not currently available in any system in our region, where the response times are very short.

## PD5. Physics-based Probabilistic Seismic Hazard Assessment (PSHA)

| PD5 | Physics-based Probabilistic Seismic Hazard Assessment |
|---|---|
| Leader | Alice-Agnes Gabriel (LMU) |
| Participants | • Bo Li, Sara Aniko Wirp, Thomas Ulrich, Fabian Kutschera (LMU)<br>• Benedikt Halldorsson, Claudia Abril (IMO)<br>• Milad Kowsari, Farnaz Bayat (UICE)<br>• Otilio Rojas, Juan E. Rodríguez, Marisol Monterrubio-Velasco, Josep de la Puente (BSC)<br>• Emanuele Casarotti (INGV)<br>• Michael Bader, Lukas Krenz, Leonhard Rannabauer (TUM) |
| Workflow | Ph-PVHA workflow; Cybershake |
| Engines | SeisSol, Exahype, AWP-ODC, SPECFEM3D |
| TRL initial | 4 |
| TRL target | 6-7 |
| TRL achieved | 6 (*) |

**(*)** Ready to raise to TRL 7-8 (Individual components of the prototype service demonstrated in operational environments). It will be tested in the exercise planned in November 2021.

### HPC Products (available software and workflows)

The two complementary workflows we developed for PD5 include:

**Workflow #1: Complex, dynamic earthquake modeling for physics-based PSHA using SeisSol**

HPC products include observationally constrained and verified 3D dynamic rupture simulations using **SeisSol** (www.seissol.org; https://github.com/SeisSol/SeisSol), a post-processing python toolbox (https://github.com/SeisSol/SeisSol/tree/master/postprocessing/science/GroundMotionParametersMaps), utilizing routines from the GMPE Strong Motion Modeller's Toolkit (GMPE-SMTK, https://github.com/GEMScienceTools/gmpe-smtk), to calculate the ground motion, and the estimation of annual seismic rates using **SHERIFS** (https://github.com/tomchartier/SHERIFS). For hazard aggregation, we combine the rupture probability output from SHERIFS and post-processed ground motion outputs from dynamic earthquake scenarios with SeisSol into fully physics-based PSHA. We note that we do not rely on empirical ground motion prediction equations (GMPE's) in this workflow. All software is open source and freely available to the community. In addition to producing physics-based PSHA, this workflow can also yield **credible worst**

**case scenarios**, and generate physics-based ground motion models (**GMMs**) that can be readily used as input to well-established hazard analysis tools to generate hazard curves or maps, such as the **OpenQuake Engine** (https://github.com/gem/oq-engine) or **OpenSHA** (http://www.opensha.org/), to design hybrid - computational and data-driven - ground motion models which are especially useful for data sparse regions.

**Workflow #2: Physics-based Probabilistic Seismic Hazard Analysis using Cybershake**

The goal of using 3D physical simulations to produce PSHA faces the challenge of the amount of events included in a probabilistic forecast. In order to partially mitigate this issue, Cybershake (https://strike.scec.org/scecpedia/CyberShake) has been designed to benefit from the reciprocity principle in elastodynamics. The strategy of the workflow is based upon selecting sites where intensity measurements are requested and recording the full stress Green tensor (SGT) solution at each fault segment that could host a future earthquake. The resulting Green tensor, computed with AWP-ODC-SGT at present time, can later be post-processed to accumulate the impact of any number of events stemming from all fault segments involved (i.e. by means of DirectSynth https://strike.scec.org/scecpedia/DirectSynth). As a result PSHA maps and intensity vs distance curves can be obtained at each site. The original software uses third-party workflow managers to orchestrate the SGT, post-processing and stochastic components for all sites and a given rupture forecast. Within ChEESE we have developed a workflow manager that can be used as an alternative. Important aspects of Cybershake relevant to PSHA is that it takes into account directivity and basin effects into the hazard maps, in comparison with GMPE-based hazard studies. Curves are generated using OpenSHA (see workflow #1 above). We remark that Cybershake uses stochastic kinematic models to represent fault rupture and assumes a flat topography at present time.

**Use case #1. Húsavík–Flatey fault zone (HFFZ), Northern Iceland**

***Summary of technological achievements***

A **single dynamic rupture earthquake scenario** for the HFFZ requires a ~3 million element model and **~3 hours computational time using 960 cores on SuperMUC-NG** at the Leibniz Supercomputing Center for one simulation, with basis functions of polynomial order 4 (5th order accurate in space and time). To assess mechanical viability, define physically limited worst case scenarios and explore a range of physically plausible initial conditions, we performed **a few hundreds of simulations** varying several representations of natural complexity, specifically fault geometrical complexity. We develop a necessary disaggregation procedure, since all earthquake scenario characteristics are not prescribed but evolve in each simulation based on first-order physical principles. To automatically detect scenarios which spontaneously develop into large, hazardous events and fulfill predefined criteria (specifically in terms of moment magnitude) for consideration for PSHA, we developed post processing python tools to automatically visualize the outputs, and calculate ground motion characteristics.

***Scientific achievements***

To perform physics-based PSHA in the Húsavík–Flatey fault Zone (HFFZ), we use SeisSol to run 3-D spontaneous dynamic rupture simulations and investigate physics-based ground motion synthetics in Northern Iceland. We construct geometric fault models of varying complexity constrained by geological and seismic data, and account for various complexities that affect the ground shaking, such as the 3-D subsurface structure, bathymetry and topography of the area, viscoelastic attenuation and off-fault plastic deformation of the host rock matrix.

We here use recently developed efficient tools and physical frameworks using a multitude of geophysical observations to define initial conditions and to verify single earthquake observables in 3D dynamic rupture (and linked tsunami simulations) as well as longer term fault slip models, which are published with international collaborators including INGV, Utrecht University, the University of Michigan and University of Oregon, USA, KAUST University Saudi Arabia, GNS New Zealand, and Kyoto University, Japan, in Palgunadi et al., BSSA 2020; Perez-Silva et al., GRL 2021; Tinti et al., EPSL 2021; Ramos et al., JGR 2021.
We also utilize open-source workflows and detailed sensitivity analysis of generic dynamic earthquake scenarios to initial conditions, published in Madden et al., GJI 2021 and Wirp et al., Frontiers, 2021. Our production scenarios can reach sustained petascale performance on various supercomputers world-wide (Krenz et al., SC21).

In addition to reproducing historic large magnitude rupture scenarios, we also vary the geometry models, hypocenter locations, and initial stress conditions to simulate mechanically possible rupture scenarios in the HFFZ. Such variations in fault geometry and stress conditions allow us to generate rupture scenarios with various magnitudes between Mw 5 and 7.3. Our simulation results, based on a large number of dynamic rupture scenarios, show that the fault geometry has a strong effect on multi-fault rupture dynamics across the HFFZ, slip levels and distribution patterns, and the final magnitude of the rupture scenarios in the HFFZ. The complex geometry model, with its 55 fault segments, separated by a variety of gaps and step overs, overall does not favor rupture scenarios that result in earthquakes larger than Mw 7. While the largest magnitude, ~Mw 7.3, can be achieved with simple fault geometry where 4 fault segments are smoothly connected. It also results in the strongest ground motion ~1.55 g in the Húsavík town. The fault roughness produces smaller magnitude ruptures and generates more high frequency ground motions. All simulated scenarios show non-linear coupled source, path and site effects on the ground motion, and yield heterogeneous ground shaking distributions along and across the fault. We observe ground shaking amplification from rupture directivity, from localized geometric complexities, such as fault gaps and kinks, and both amplification and shielding from topography. The attenuation of the physics-based ground motion with distance from the faults shows magnitude consistent attenuation relationship, especially in the near-fault region, and overall good agreement with empirical ground motion models (GMMs) specific to the Southern Iceland Seismic Zone, a tectonically and seismically symmetric zone with

Northern Iceland. The ground motion variability, usually a constant value in GMMs, changes with distance to the fault, and has higher values for unilateral than bilateral rupture scenarios. This suggests the GMMs are more complex than the empirical ones derived from limited data, and that complexities in the rupture models are needed to improve the ground motion simulation and hazard analysis as well as the uncertainty quantification.

**Use case #2. CyberShake application at the South Iceland Seismic Zone (SISZ) and the Reykjanes Peninsula Oblique Rift (RPOR) regions**

***Summary of technological achievements***

A main technological step for the CyberShake (CS) application in South Iceland that may facilitate its migration to other regions is the implementation of a workflow manager called *UnifiedCSWFlow* that runs on Marenostrum 4 (MN4).

The original CS installation performed by SCEC on MN4 corresponds to a version capable of computing hazard curves for a single site. In previous executions in California, the developers relied on external tools for workflow management (i.e. Pegasus WMS or HTCondor) which are not property of SCEC and hence not directly available to ChEESE. The orchestration of CS runs, within ChEESE, is performed now with *UnifiedCSWFlow,* a Python (> 3.6 compliant) code that runs on MN4 interactive nodes and prepares and submits orderly all needed steps of CS, thus reducing the tedious manual interaction of previous CS exercises. *UnifiedCSWFlow* interacts with the non-relational CS database by means of a simple spreadsheet that collects all information provided by the user, such as, e.g. the geological model, fault geometry or Earthquake Forecast Model (EFM). *UnifiedCSWFlow* handles all script preparation, job submission and result merging automatically. It is currently working on MN4 and tested for Southern Iceland but its capabilities could be easily expanded to other regions or HPC systems. It has been one of the main drivers of a demonstration exercise for PD5 users within ChEESE.

***Scientific achievements***

Our main contribution is the validation of CS simulations of intensity measures in the RPOR-SISZ region. This is an important step prior to hazard studies, where the intensity measures are compared to existing solutions (i.e. GMMs) in order to find the validity of the geological model and the modelling engine (AWP, in this case). This test considers a fault model of 36 planar near-vertical North-South oriented dextral transform faults. Each of them is forced to rupture at three moment magnitudes, including the 16- 50- and 84-percentile probable maximum magnitude of each fault. These Mw values are chosen according to the Mai & Beroza (2000) effective source-area empirical relation that is assumed valid for earthquakes in this region. Our findings, comparing spectral acceleration up to 1 Hz and for three sites, result in good agreement between GMMs and simulation results, mostly within one standard deviation. This includes events with 5.2<Mw<7.2 and at maximum distances of ~100 km. We remark that given the different assumptions in GMMs and physical simulations (e.g. directivity) a complete agreement is not expected.

In our simulations, we have further compared three velocity models for this region, namely the 1D South Iceland Lowland (SIL) model, the 1D RPOR-SISZ model (a combination of two regional 1D models), and a 3D tomographic SISZ-RPOR model, all of them provided by IMO. The impact of the models in the results does not appear to be very high, and we agree upon relying on the 1D RPOR-SISZ for further runs.

A further question was posed by IMO, as the obtained stochastic fault rupture models generated by the CS's Graves-Pitarka genv3 displayed too few large scale features, as opposed to what would be expected at SISZ faults for large events. In order to obtain more qualitatively realistic rupture scenarios we have introduced genv5 and improved its parameterization. Consequently, slip distributions obtained are more realistic although the impact in intensity measurements is negligible. We will thus keep genv5 as preferred stochastic slip generator in further exercises.

The validation exercise has thus allowed us to answer fundamental questions regarding the calibration and specialization of CS to Southern Iceland (i.e. geological model, rupture generator), prior to attaining PSHA maps. Our next, immediate, step will be validating CS for smaller earthquakes, which may originate away from the main faults used in the current validation.

**PD5 Scientific Products**

***Published papers***


- Madden, Elizabeth, et al. (2021), Linked 3D modeling of megathrust earthquake-tsunami events: from subduction to tsunami run up, Geophysical Journal International, 224(1), 487–516, doi:doi:10.1093/gji/ggaa484.
- Ramos, Marlon, Yihe Huang, Thomas Ulrich, Duo Li, Alice-Agnes Gabriel, and Amanda Thomas (2021), Assessing Margin-Wide Rupture Behaviors along the Cascadia Megathrust with 3-D Dynamic Rupture Simulations, Journal of Geophysical Research - Solid Earth, 126, doi:10.1029/2021JB022005.
- Tinti, Elisa, Emanuele Casarotti, Thomas Ulrich, Duo Li, Taufiqurrahman Taufiqurrahman, and Alice-Agnes Gabriel (2021), Constraining families of dynamic models using geological, geodetic and strong ground motion data: the Mw 6.5, October 30th, 2016, Norcia earthquake, Italy, EPSL, 576, 117237, doi:10.1016/j.epsl.2021.117237.
- Wirp, Sara Aniko, Alice-Agnes Gabriel, Elizabeth H Madden, Maximilian Schmeller, Iris van Zelst, Lukas Krenz, Ylona van Dinther, and Leonhard Rannabauer (2021), 3D linked subduction, dynamic rupture, tsunami and inundation modeling: dynamic effects of supershear and tsunami earthquakes, hypocenter location and shallow fault slip, Frontiers in Earth Science, Geohazards and Georisks, doi:10.3389/feart.2021.626844.

- Perez-Silva, Andrea, Duo Li, Alice-Agnes Gabriel, and Yoshihiro Kaneko (2021), 3D modeling of long-term slow slip events along the flat-slab segment in the Guerrero Seismic Gap, Mexico, Geophysical Research Letters, 48(13), doi:10.1029/2021GL092968.
- Palgunadi, Kadek Hendrawan, Alice-Agnes Gabriel, Thomas Ulrich, José Ángel Lopéz-Comino, and Paul Martin Mai (2020), Dynamic Fault Interaction during a Fluid‑Injection‑Induced Earthquake: The 2017 Mw 5.5 Pohang Event, Bulletin of the Seismological Society of America, 110(5), 2328–2349, doi:10.1785/0120200106

***Conferences, seminars***

- Li, B., Gabriel, A. A., Ulrich, T., Abril Lopez, C., Halldorsson, B., & Bader, M. (2020, December). Dynamic earthquake rupture scenarios and physics-based seismic hazard assessment for the segmented Húsavík-Flatey fault zone, North Iceland. In AGU Fall Meeting Abstracts (Vol. 2020, pp. S036-0001).
- Li, B., Gabriel, A. A., Ulrich, T., Abril Lopez, C., Halldorsson, B., & Bader, M. (2021, December). Physics-based rupture models and ground shaking simulations in the Húsavík–Flatey fault zone, Northern Iceland. In AGU Fall Meeting Abstracts.
- Abril, C., Mai, M., Halldórsson, B., Li, B., Gabriel, A., and Ulrich, T.: Ground motion simulations for finite-fault earthquake scenarios on the Húsavík-Flatey Fault, North Iceland, EGU General Assembly 2021, online, 19–30 Apr 2021, EGU21-15557, https://doi.org/10.5194/egusphere-egu21-15557, 2021.
- Kutschera, F., Wirp, S. A., Li, B., Gabriel, A.-A., Halldórsson, B., Abril, C., and Rannabauer, L.: Linking dynamic earthquake rupture to tsunami modeling for the Húsavík-Flatey transform fault system in North Iceland, EGU General Assembly 2021, online, 19–30 Apr 2021, EGU21-15891, https://doi.org/10.5194/egusphere-egu21-15891, 2021.
- Krenz, Lukas, Carsten Uphoff, Thomas Ulrich, Alice-Agnes Gabriel, Lauren Abrahams, Eric Dunham, and Michael Bader (2021), 3D Acoustic-Elastic Coupling with Gravity: The Dynamics of the 2018 Palu, Sulawesi Earthquake and Tsunami, SC'21: Proceedings of the international conference for high performance computing, networking, storage and analysis, in press.
- Rojas, O., Rodriguez, J. E., de la Puente, J., Callaghan, S., Abril, C., Halldorsson, B., Li, B., Gabriel, A. A., and Olsen, K.: Towards physics-based PSHA using CyberShake in the South Iceland Seismic Zone, EGU General Assembly 2021, online, 19–30 Apr 2021, EGU21-7880. (https://doi.org/10.5194/egusphere-egu21-7880, 2021).
- Rojas, O., Rodriguez, J. E., de la Puente, J., Callaghan, S., Abril, C., Halldorsson, B., Li, B., Gabriel, A. A., and Olsen, K.: Insights on Physics-based Probabilistic Seismic Hazard Analysis in South Iceland using CyberShake. AGU Fall Meeting 2021.

- C. Abril, B. Halldórsson, P.M. Mai and S. Jonsson (2021) Ground-motion simulations for finite-fault earthquake scenarios on the Reykjanes Peninsula, South Iceland. 37th General Assembly of the European Seismological Commission, 19-24 September 2021, Athens, Greece. Paper No. 593.
- Benedikt Halldorsson, Farnaz Bayat, Milad Kowsari, Claudia Abril (2021). New 3D fault system models of the two transform zones of Iceland for physics-based seismic hazard assessment. 52nd Nordic Seismology Seminar, Reykjavík, Iceland, 18-20 October 2021.
- Farnaz Bayat, Benedikt Halldorsson, Milad Kowsari (2021). On the calibration of the first 3D transform fault system model of the South Iceland Seismic Zone and Reykjanes Peninsula Oblique Rift. 52nd Nordic Seismology Seminar, Reykjavík, Iceland, 18-20 October 2021.
- Milad Kowsari and Benedikt Halldorsson (2021). Sensitivity of the Seismic Hazard Maps to the Selected Ground Motion Models: A Case Study of North Iceland. 52nd Nordic Seismology Seminar, Reykjavík, Iceland, 18-20 October 2021.
- Kowsari, Halldorsson, B., Kijko, A., Bessason, B., & Jónasson, K. (2021). Comparison of Seismicity Parameters from Various Earthquake Catalogues for Iceland. 17th World Conference on Earthquake on Earthquake Engineering (17WCEE), Paper No. 3837.

**Validation.**

| **Functional requirement**<br><br>**(these are examples)** | **Target (from D5.1)** | **Achieved** | **Validated (YES/NO)** |
|---|---|---|---|
| Time to solution. | Not a constraint | 3500 s (Cybershake for 1 site, including AWPX; AWPY and DirectSynth with 36 nodes)<br><br>~3 hours using 940 cores of SuperMUC-NG (HFFZ, SeisSol) | not applicable |
| Resolution. | Resolution of ground motion in the engineering frequency band | 1 Hz (Cybershake, with 100m grid)<br><br>1 Hz, and 3 Hz in the near fault region (HFFZ, SeisSol) | YES |
| Number of simulations. | Large ensemble size required for extracting probabilistic hazard information.<br>Hundreds/Thousands of forward simulations | 3 sites, each running a 2 reciprocal SGTs (Cybershake)<br><br>1-2 hundred (HFFZ,SeisSol) | YES (lower bound) |
| Data formats. | Asynchronous output, local time stepping enabled | PSA files, Cybershake format<br><br>hdf5 (SeisSol) | YES |

Table 5.1. Validation criteria for PD5.

**Involvement of end-users**

The Icelandic Meteorological Office (IMO) is an end-user of the product of PD5, and the sensitivity analysis of a physics-based PSHA in a transform zone of Iceland is of direct relevance. For this purpose the focus has been on the region of highest seismic risk in Iceland, the South Iceland Seismic Zone and the Reykjanes Peninsula Oblique Rift. As well as acting as end-user, the IMO is a participant of the ChEESE project with the specific purpose of defining the center, body and range of realistically possible values of the physical parameters required for Physics Based PSHA (PB-PSHA). That includes the definition of the confinement of seismic sources in the SISZ-RPOR, in the form of a complete 3D finite-fault fault system model for the region. In addition, such a model has also been defined for the other transform zone in the country, the Tjörnes Fracture Zone of North Iceland. The models are complete in terms of the centre estimates of most likely fault locations, fault dimensions, and their seismic activity manifested by each fault's long-term slip-rate.

The efforts towards a PB-PSHA in the two transform zones are carried out in parallel but with two different approaches and using different tools. In this manner, the efforts in the SISZ-RPOR are focused on the probabilistic approach using the CyberShake simulation platform, while those in the TFZ are focused on a deterministic approach using SeisSol and SORD earthquake rupture simulators. To elaborate, the efforts in the SISZ-RPOR are focused on the Monte Carlo simulation of earthquake scenarios that fulfill the slip-rates defined by the 3D fault system model. For the TFZ, likely worst-case earthquake scenarios are postulated on known fault systems close to urban areas. In both cases this involves the simulation of strong-motions at hypothetical grid-locations in the macroseismic region and the affected surrounding areas. Together, these two complementary approaches provide a more complete view of the sensitivity of the PB-PSHA, which the IMO as end-user can evaluate and compare with the state-of-the-art knowledge of seismic sources and seismic activity in the transform zones. The comparison among other things involves contrasting the PB-PSHA results with those of standard engineering approach to PSHA based on empirical ground motion attenuation relations, in addition to comparing to existing hazard maps, including the latest harmonized European efforts in that regard, the European Seismic Hazard Model of 2020 (ESHM20).

**Impact**

The current state-of-the-art of PSHA based on the engineering approach is primitive in the sense that it uses:

(1) magnitude-frequency relationships (i.e., Gutenberg-Richter) for simple seismic area sources on the basis of a statistical analysis of historical earthquake catalogues that are riddled with uncertainties in both magnitudes and locations.
(2) empirical functions for the scaling of ground motion peak-parameters for which the associated variability has not been shown to reduce significantly over the last

two decades, despite vast new data. Moreover, the uncertainties are estimated to be largely associated with the complexity of earthquake rupture, which is not modeled.

The uncertain variability of the seismic activity, along with the variability of ground motion models cascades to greatly increase the uncertainty associated with the PSHA results. This in turn counteracts the effective and optimal seismic design. However, in PD5, towards PB-PSHA for the transform zones of Iceland, we have comprehensively addressed these limitations:

(1) Developed new 3D earthquake source models that are calibrated to the rate of tectonic motions and the geometry of the fault system. It is thus developed from fundamental assumptions of the physics of the earthquake sources, and fully specifies the long-term seismic activity of the system. The activity of the new 3D fault models produces synthetic earthquake catalogues that are in agreement with the historical and instrumental earthquake catalogues. Moreover, the 3D models allow the subzonation of the fracture zones in terms of robust physics-based estimates of their seismic activity, something that was not possible before due to the scarce catalogues for subzones.

(2) Modeled the heterogeneous distribution of earthquake source processes on realistic fault rupture scenarios, modeled using the dynamic and kinematic modeling approaches, for multiple realizations thus effectively investigating the sensitivity of the ground motion amplitudes to variations in earthquake rupture.

Thus, our dynamic rupture simulations provide an estimation of the physically plausible worst scenario and the ground shaking in HFFZ. The heterogeneous distribution of physics-based ground motion suggests the ground shaking pattern is able to account for much more source complexity than current GMMs. Our results also demonstrate the potential to establish physics-based GMMs to complement existing GMMs, and apply for probabilistic seismic hazard assessment, specifically in data-limited regions.

In the case of CyberShake the platform installation in MN4 and applications to SISZ-RPOR is by itself an achievement beyond ChEESE objectives, and paves the way for extending its applications to other seismic regions of Europe, with Iceland being the first application outside California. CS ground motions estimations in SISZ-RPOR, although limited to kinematic rupture parameterization and flat topographies, may serve as a reference at low frequencies ≤1Hz, for results modeled with dynamic rupturing.

CyberShake advantage relies on the parallel computation of a large suite of earthquake simulations given by an ERF to allow physics-based PSHA.

The challenge in engineering seismology and its application to PSHA is to incorporate PB-PSHA methods in a consistent manner, and then to extend its validity and thus its application to higher frequencies, up to 10 Hz. This is a challenge, but it cannot be achieved in a physically robust manner unless starting with a reliable and consistent earthquake rupture model that is comprehensive and efficient. Comprehensive in the sense that it is consistent with state-of-the-art dynamic rupture models, but efficient in the sense that an “equivalent” kinematic rupture model is applied that is much more computationally efficient. Then, small-scale heterogeneities in the rupture process will be added in order to simulate higher frequency motions in a physically consistent manner from the same earthquake source model.

| use-case # | Architecture | Number of CPU/GPU per application | Number of applications | maximum achieved FLOPS | Total used CPU/GPU hours |
|---|---|---|---|---|---|
| **Cybershake 1 Hz demo** | Intel Skylake CPUs | 1728 cores | 3 | 11.5 GFLOPS | 302,400 core-hours |
| **SeisSol** | Intel Skylake CPUs | 960 cores (20 nodes)<br>384 cores (8 nodes) | 100-200<br>1 | ~12850 GFLOPS*<br>~8880 GFLOPS** | ~420,000 core-hours<br>- |
| | DGX A100 (AMD EPYC 7742) | 8x A100 GPUs with 128 CPU cores (1 node) | 1 | ~7515 GFLOPS** | - |

Table 5.2. For SeisSol user case, there are three rows: 1. - performance obtained with convergence order 5 and double precision (DP) floating point format (*); 2. and 3. - performance obtained with convergence order 6 and double precision (DP) floating point format (**);

## PD6. Probabilistic Volcanic Hazard Assessment (PVHA)

| **PD6** | **Probabilistic Volcanic Hazard Assessment** |
|---|---|
| **Leader** | Laura Sandri |
| **Participants** | • Beatriz Martinez Montesinos, Antonio Costa, Giovanni Macedonio (INGV)<br>• Manuel Titos, Sara Barsotti (IMO)<br>• Arnau Folch, Leonardo Mingari (CSIC) |
| **Workflow** | ChEESE PVHA workflow |
| **Engine** | Fall3D |
| **TRL initial** | 3 |
| **TRL target** | 6 |
| **TRL achieved** | 6 (*) |

**(*)** Ready to raise to TRL 7-8 (Individual components of the prototype service demonstrated in operational environments). It will be tested in the exercise planned on 4th November 2021.

**HPC Products (available software and workflows)**

Taking as a starting point a previously developed prototype tool, we have built an HPC-based workflow to perform a probabilistic assessment of airborne ash and ashfall related volcanic hazard both in the short and long term. For that we take advantage of the new FALL3D flagship code, the performance and productivity optimization (POP) and the workflow management services provided by WP3, as well as the computational **resources awarded by PRACE at Joliot-Curie at TGCC-CEA (France)**. As a result, hazard and probability maps for ground load, ash concentration at different Flight Levels (FL), arrival times for specific ash concentration (again at different FL) and persistence maps of specific ash concentrations (again at different FL) have been produced and reported.

Alongside this pilot demonstrator, a 2 staged workflow has been developed to generate and process the data. The first stage, samples and generates the eruptive scenarios, while the second stage post-processes the results and gets the hazard maps. Since the computational cost of the first stage is low, the optimization has been carried out in the second stage of the workflow (post-processing).

Furthermore, given the high number of simulated scenarios, a module to compress the data without significant information losses to obtain the risk maps has also been implemented in this pilot demonstrator. This allows us to move the data outside the HPC environment where they were obtained.

**Use case #1. Campi Flegrei Long-Term Probabilistic Volcanic Hazard Assessment (CF-LTPVHA)**

***Summary of technological achievements (HPC performance, etc)***

For this use case we have run **4500 FALL3D simulations** (1500 per 3 eruption size classes) on a grid 2000km x 2000km, 2km resolution, on Joliot-Curie at TGCC-CEA (France), within a PRACE-funded project in association with PD3 and PD12.

We have used the performance and productivity optimization (POP) service provided by ChEESE to optimize the parallelization of the part of the code in charge of analyzing the FALL3D simulations and calculating the necessary probabilities to carry out the assessment. As a result we have decreased execution time by several orders of magnitude. In addition, within the ChEESE WP3 workflow management, a tool has been developed capable of executing the different workflow modules on different servers and managing the necessary data flow.

Regarding the HPC environment, we ran a few preliminary cases (with grid size similar to that of the PVHA, i.e. 50M grid points and 12 particle bins) changing the configuration of nodes and cores used to optimize the energy consumption and computing time. For our grid size, parallel efficiencies already drop to 70% with only 1036 processors (32 nodes). Then, considering the resolution of our domain (0.025°), and the total grid points 35M (1040*920*35), we fixed the number of nodes to 16 and the number of cores to 768. This configuration allows decomposing the grid points into 32*24*1 (X,Y ,Z) subdomains of more than 30 points per spatial dimension. As a result, we increased the speed-up 16 times and the parallel efficiency was about 90%.

***Scientific achievements***

With this use case we have demonstrated the feasibility of providing long-term probabilistic hazard assessment from tephra ground load and airborne ash concentration. With “long-term” we mean that the probability of the hazard impact is computed on a time window of years to decades, and the assessment is mostly based on data from geological record for what concerns volcanic recurrences, and simulations based on a statistic of wind profiles for the tephra dispersal. The assessment here has been made:

- on a large-scale (2000km x 2000km) and high-resolution (2km x 2km) domain,
- exploring the variability in the Eruption Source Parameters linked to the different possible eruptive sizes
- exploring the variability in wind conditions.

This is feasible thanks to HPC resources, and it represents an improvement over the current activity at Osservatorio Vesuviano, in which three fixed eruptive scenario (corresponding to 3 reference sizes, each with fixed Eruption Source Parameters) are simulated over a 900 km x 900 km grid, only for ground load purposes.

PD6 has produced and validated maps for airborne ash concentrations (in terms of maximum concentration, concentration persistence and arrival time above a given concentration) and of ground load of tephra over the target domain, considering a continuous spectrum of eruptive size classes (and related Eruption Source Parameters) and accounting for 20 years of wind variability.

***Scientific Products (including publications)***

Hazard and probability maps for ground load, ash concentration at different Flight Levels (FL), arrival times for specific ash concentration (again at different FL) and persistence maps of specific ash concentrations (again at different FL) have been produced and reported in D4.6, to demonstrate the use case.

A scientific publication for Frontiers in Earth Science is in preparation.

***Conferences, seminars***

- Martínez Montesinos, B., Titos, M., Sandri, L., Barsotti, S., Macedonio, G., and Costa, A.: Probabilistic Tephra Hazard Assessment of Campi Flegrei, Italy, EGU General Assembly 2021, online, 19–30 Apr 2021, EGU21-7595, https://doi.org/10.5194/egusphere-egu21-7595, 2021.

**Use case #2. Campi Flegrei Short-Term Probabilistic Volcanic Hazard Assessment (CF-STPVHA)**

***Summary of technological achievements (HPC performance, etc)***

For this use case we have run **180 FALL3D simulations for each eruptive size class** (Large, Medium, Small) at Campi Flegrei on a grid 2000km x 2000km, 2km resolution, on Joliot-Curie at TGCC-CEA (France), within a PRACE-funded project in association with PD3 and PD12. We have performed this process for days 5-6-7th December, 2019 carrying out a total of 1620 simulations.

All the simulations have been run with meteorological data for the 5-6-7th December, 2019.

As in the previous case, POP and WP3 workflow management services have been used allowing communication with the server responsible for collecting the monitoring data of Campi Flegrei (Vesuvius observatory) and a faster development of the short-term assessment.

Regarding HPC environment, we used the configuration chosen for Use case #1.

***Scientific achievements***

With this use case we have demonstrated the feasibility of providing short-term probabilistic hazard assessment from tephra ground load and airborne ash concentration. With “short-term” we mean that the probability of the hazard impact is computed on a time window of days to weeks, and the assessment is mostly based on data from the monitoring system for what concerns volcanic occurrences, and simulations based on the most updated wind forecast for the tephra dispersal. The assessment here has been made:

- on a large-scale (2000km x 2000km) and high-resolution (2km x 2km) domain
- exploring the variability in the Eruption Source Parameters linked to the different possible eruptive sizes.

This is all feasible thanks to HPC resources, and it represents an improvement over the current activity at Osservatorio Vesuviano, in which every day only three fixed eruptive scenarios (corresponding to 3 reference sizes, each with fixed Eruption Source Parameters) is simulated over a 200 km x 200 km grid.

PD6 has produced and validated maps for airborne ash concentrations (in terms of maximum concentration, concentration persistence and arrival time above a given concentration) and of ground load of tephra over the target domain, considering a continuous spectrum of eruptive size classes (and related Eruption Source Parameters) for a specific time period (5 to 7 December 2019).

***Scientific Products (including publications)***

Maps for a specific time period (5-7 December 2019), to demonstrate the use case, have been produced and reported in D4.6.

A scientific publication for Frontiers in Earth Science is in preparation.

A further validation in an operational environment of the workflow in this use case will be carried out in the simulation exercise that will be held on 4th November 2021.

***Conferences, seminars***

- Martínez Montesinos, B., Titos, M., Sandri, L., Barsotti, S., Macedonio, G., and Costa, A.: Probabilistic Tephra Hazard Assessment of Campi Flegrei, Italy, EGU General Assembly 2021, online, 19–30 Apr 2021, EGU21-7595, https://doi.org/10.5194/egusphere-egu21-7595, 2021.

**Use case #3. Jan Mayen Long-Term Probabilistic Volcanic Hazard Assessment (JM-LTPVHA)**

***Summary of technological achievements (HPC performance, etc)***

For this use case we have run 1500 FALL3D simulations for the Large eruptive size class at Jan Mayen, on a grid 2000km x 2000km, 2km resolution, on Joliot-Curie at TGCC-CEA (France), within a PRACE-funded project in association with PD3 and PD12.

Also, on the same grid and with the same HPC resources, we have run 1500 FALL3D simulations for the Medium size class, here implementing a novel strategy to simulate tephra dispersal from long-lasting and pulsating eruption of given total erupted mass.

As in the previous cases, POP and WP3 workflow management services have been used.

Regarding HPC environment, we used the configuration chosen for Use case #1.

***Scientific achievements***

With this use case we have demonstrated the feasibility of providing long-term (years to centuries) probabilistic hazard assessment from airborne ash concentration at Jan Mayne, a little studied volcanic island located in the far-North Atlantic Ocean under some major air-routes. In particular, we have quantified hazard:

- on a large-scale (2000km x 2000km) and high-resolution (2km x 2km) domain,
- exploring the variability in the Eruption Source Parameters linked to the different possible eruptive sizes
- exploring the variability in wind conditions.

This is all feasible thanks to HPC resources, and it represents a major improvement over the state-of-the-art for this volcano, which has been so far little studied.

We have also developed a novel strategy to explore the Eruption Source Parameter variability for pulsating events (typical of the Medium size class at Jan Mayen), ejecting tephra in the atmosphere through pulses lasting a few hours to days but clustering over a period of a month. This is especially time-consuming to simulate. In PD6 we developed a strategy to sample the Eruption Source Parameters in different pulses within ranges coming from literature and consistent with the total erupted mass in the whole eruption.

PD6 has produced and validated maps for airborne ash concentrations (in terms of maximum concentration, concentration persistence and arrival time above a given concentration, at different FL) considering a continuous spectrum of eruption source parameters and accounting for 20 years of wind variability.

***Scientific Products (including publications)***

Hazard and probability maps for ash concentration at different FL, arrival times for specific ash concentration (again at different FL) and persistence maps of specific ash concentrations (again at different FL) have been produced and reported in D4.6, to demonstrate the use case.

***Papers***

- Titos, M., Martínez Montesinos, B., Barsotti, S., Sandri, L., Folch, A., Mingari, L., Macedonio, G., and Costa, A.: Assessing potential impact of explosive volcanic eruptions from Jan Mayen Island (Norway) on aviation in the North Atlantic, Nat. Hazards Earth Syst. Sci. Discuss. [preprint], https://doi.org/10.5194/nhess-2021-264, in review, 2021.

***Conferences, seminars***

- Titos, M., Martínez, B., Barsotti, S., Sandri, L., Folch, A., Mingari, L., Costa, A., and Macedonio, G.: Assessing potential impacts on the air traffic routes due to an ash-producing eruption on Jan Mayen Island (Norway), EGU General Assembly 2021, online, 19–30 Apr 2021, EGU21-7594, https://doi.org/10.5194/egusphere-egu21-7594, 2021.

**Validation.**

| **Functional requirement**<br><br>**(these are examples)** | **Target (from D5.1)** | **Achieved** | **Validated (YES/NO)** |
|---|---|---|---|
| Domain size and resolution | Run FALL3D on a large-scale (2000km x 2000km) and high-resolution (2km x 2km) domain | 2000km x 2000km large domain at 2km x 2km resolution | YES |
| Number of simulations. | 1000 FALL3D simulations per eruptive size in Long-Term use cases | 1500 per eruptive size | YES |
| Number of simulations | 100 FALL3D simulations per eruptive size in Short-Term use cases | 180 per eruption size, per each day of the explored period explored | YES |
| Data formats | Defined in task 4.3 as NetCDF | hazard output formatted accordingly | YES |
| Time for simulations | FALL3D simulations on the large scale domain in less than 3 hours after each service launch (for Short-Term use cases) | We are validating it in the exercise (D5.4) to be held on 4th November | |
| Time for workflow | Workflow runtime less than 1.5 hours after each service launch (for Short-Term use cases) | We are validating it in the exercise (D5.4) to be held on 4th November | |
| Accessibility to data for workflow input | Real-time access to Volcano Observatory database of monitoring data (for Short-Term use cases) | Osservatorio Vesuviano only | YES |
| Accessibility to data for workflow input | Access to ERA5 reanalysis dataset (for Long-Term use cases) | ERA5 reanalysis dataset integrated | YES |
| Accessibility to data for workflow input | Automatic and timely access to weather forecast (for Short-Term use cases) | GFS forecast accessed automatically.<br><br>We are validating it in the exercise to be held on 4th November | |

Table 6.1. Validation criteria for PD6.

**Involvement of end-users**

We have interacted with PLINIVS (who is a component of the Italian Civil Protection system) and ARISTOTLE. With PLINIVS, we have defined standard format and type of products from the Short-Term use case at Campi Flegrei, that we are validating in the exercise on 4th November 2021.

With ARISTOTLE, we are defining the hazard and probability maps of interest that we can produce in the Short-Term use case at Campi Flegrei. Again, this will be validated in the exercise on 4th November 2021.

**Impact**

The potential impact demonstrated in PD6 is the possibility to overcome the current limits in the simulation of volcanic processes (in our case volcanic ash dispersal), which is only possible due to the HPC integration and computational resources acquired. In particular, with the Tier-0 HPC integration and resources, it is possible to:

- quantify probabilistic hazard (here posed by tephra ground accumulation and by airborne ash at various flight levels) over a large-scale (thousands by thousands km) and high-resolution (of the order of 1km) target domains. This allows exploring the unlikely effects posed by low-probability but high-consequent events, especially at distal locations from the volcanic source;
- quantify the uncertainty related to the unknown Eruption Source Parameters linked to the type and scale of future eruptive events, and to wind variability, by exploring statistically significant ensembles of simulations in which the Eruption Source Parameters and wind profiles are randomly sampled from suitable probability distributions, and from statistical analysis of ensemble simulation results.

The PD has demonstrated the feasibility of such applications in different volcanic context (a highly inhabited region where ground load can be a significant risk, and at a remote volcanic island where the hazard is mainly posed on aviation routes) and over different time scales:

- short-term assessment (days to week) where the monitoring information can be used to constrain the eruption probability and the position of the vent, and where the weather forecast is a primary input, providing hazard products mainly for crisis managements, and
- long-term assessment (years to decades) where the geological and historical records can be used to evaluate the eruption probability and the position of the vent, and where the wind climate (through a statistically significant sample of wind profiles) is a primary input, providing hazard products mainly for long term planning.

| use-case # | Architecture | Number of CPU/GPU per application | Number of applications | Total used CPU/GPU hours |
|---|---|---|---|---|
| **1** | IRENE-SKYLAKE. 1656 Intel-skylake 2.7 GHz bi-processor with 24 cores per processor, a total of 79,488 computing cores and a power of 6.86 Pflop/s, 192 GB of DDR4 memory/node | 768 | 4500 runs | 768 * 15000 hours (between 1.5 and 10 hours per run, approx.) = 11.5M |
| **2** | idem | 768 | 1620 runs | 768 CPU * 5000 hours (between 1.5 and 10 hours per run, approx.) = 3.9M |
| **3** | idem | 768 | 5000 runs | 768 * 12500 hours.(between 1.3 and 3.8 hours per run) = 9.6M |

Table 6.2. Computational performance

## PD7. Probabilistic Tsunami Hazard Assessment

| PD7 | Probabilistic Tsunami Hazard Assessment (PTHA) |
|---|---|
| **Leader** | Finn Løvholt (NGI) |
| **Participants** | NGI (Steven Gibbons, Finn Løvholt, Malte Vöge, Sylfest Glimsdal)<br><br>INGV (Manuela Volpe, Stefano Lorito, Fabrizio Romano, Jacopo Selva, Roberto Tonini, Beatriz Brizuela, Angela Stallone)<br><br>UMA (Carlos Sánchez-Linares, Jorge Macías, Marc de la Asunción, Manuel J. Castro, José Manuel González)<br><br>CINECA (Piero Lanucara) |
| **Workflow** | Shell and Python script-based PTHA Workflow |
| **Codes** | Tsunami-HySEA |
| **TRL initial** | 3 |
| **TRL target** | 5-7 |
| **TRL achieved** | 6* |

**(*)** Expected to be raised to TRL 7-8 within the project duration, once integrated with WMS-light workflow manager.

**HPC Products (available software and workflows)**

The ChEESE PD7, Probabilistic Tsunami Hazard Assessment (PTHA), consists of a workflow for performing local PTHA on high resolution grids through inundation modelling. PD7 is elaborately described in ChEESE deliverable D4.7, and further the outline of the service in D5.1. Hence, we will not review additional details already provided in these past deliverables here, but limit it to reviewing the basis of the workflow and summarizing the progress made in terms of advancing the TRL and making the software operational.

**Advancement in state-of-the-art, TRL level, and capability in operational use:** The PD7 workflow advances a previous PTHA based on offshore tsunami heights to local high-resolution inundation analysis and enables this workflow to operate efficiently on a Tier-0 system resolving uncertainties adequately. Through the **PRACE project TsuHazAP**, this capability has been demonstrated by carrying out what is likely the most extensive PTHA to date globally on the **Marconi100 machine**. The number of scenarios simulated is of the order 2-3 million, which is unprecedented elsewhere in the scientific community. This dataset can serve as an international benchmark for testing PTHA codes and workflows.

The successful application on Marconi100 supports the reassessment of the TRL level to about 6-7 as **the code is ready to be used operationally**. More details are given under the test case description. Presently, the workflow is also being integrated with the workflow manager WMS light, which is expected to increase the TRL level further.

**Functionality of software:** The workflow and the connections between the different components of the overall workflow are shown in Figure 7.1, while the micro workflow consisting of the HPC part is shown in Figure 7.2. The workflow shown in Figure 7.1 is effectively executed in the following steps, where each step consists of a script that executes a given block of the overall workflow:

1. **Select the study area.** Allows the user to set up high-resolution grids used in the simulation and select the Points of Interest (PoIs) that carry the hazard information from the existing regional NEAMTHM18 model from the TSUMAPS-NEAM project. The method is also generally applicable towards other assessments if scenario parameters and rates are provided in the right format.
2. **Provide a first order hazard and source screening.** Extract the hazard curves from NEAMTHM18. Construct a branch of scenarios and metadata from the previous assessment that forms the background for the hazard analysis.
3. **Perform scenario disaggregation.** Select the scenarios from the TSUMAPS-NEAM assessment that contribute most strongly to the hazard at the location of interest. The total (mean) hazard offshore from the existing PTHA model can be reproduced at different levels of approximations, and a higher percentage (higher accuracy) implies that a larger number of scenarios needs to be executed and that more computational resources are needed.
4. **Collect scenario rates.** For the selected scenario, collect the annual rates ensembles (i.e. the rates from the family of models in NEAMTHM18 describing each scenario considering the epistemic uncertainty).
5. **Perform refinement of the scenarios.** The scenarios are refined with respect to earthquake focal mechanisms and location (crustal sources) and slip distribution (subduction sources), to better resolve source uncertainties.
6. **Perform the scenario simulations.** The HPC part of the workflow (Figure 7.2) where a high number of inundation simulations on local nested grids are carried out using the Flagship code Tsunami-HySEA.
7. **Perform the hazard aggregation.** Cumulative probabilities of specified levels of tsunami inundation are obtained by aggregating the probabilities and output for all simulations.
8. **Visualization.** Hazard maps at the local scale are produced from the hazard curves computed at the previous step.

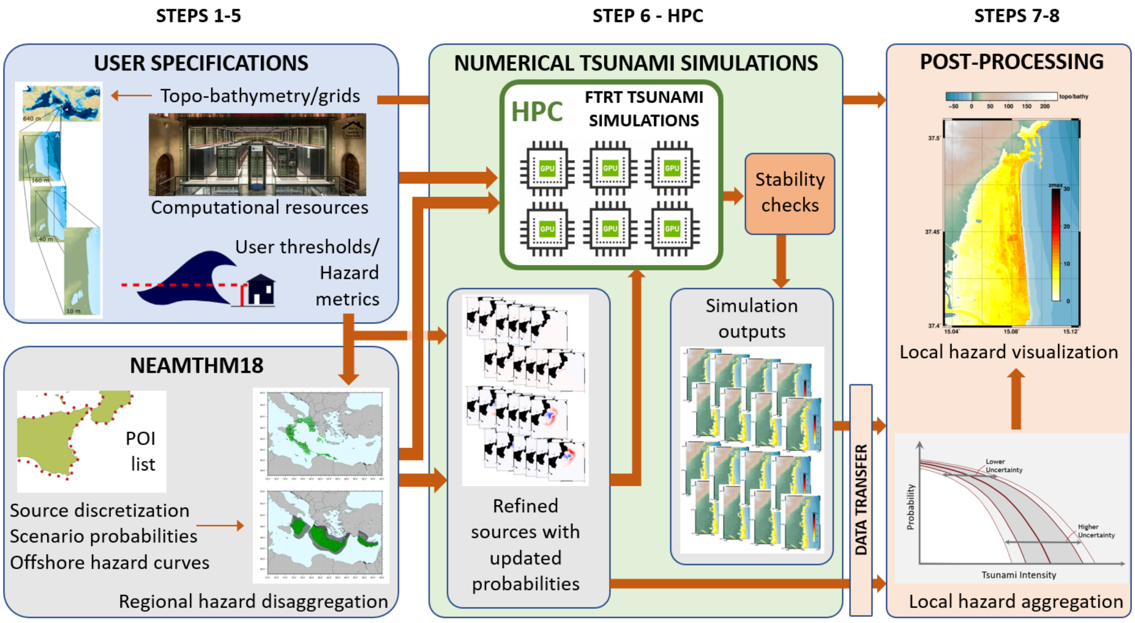


Figure 7.1. PTHA workflow designed to minimize the data transfer requirements to and from HPC resources (green box). The most substantial data transfer is indicated in which maximum height, momentum flux, and deformed bathymetry from the highest resolution grid (typically of the order 10 MB per simulation) are transferred for hazard aggregation.

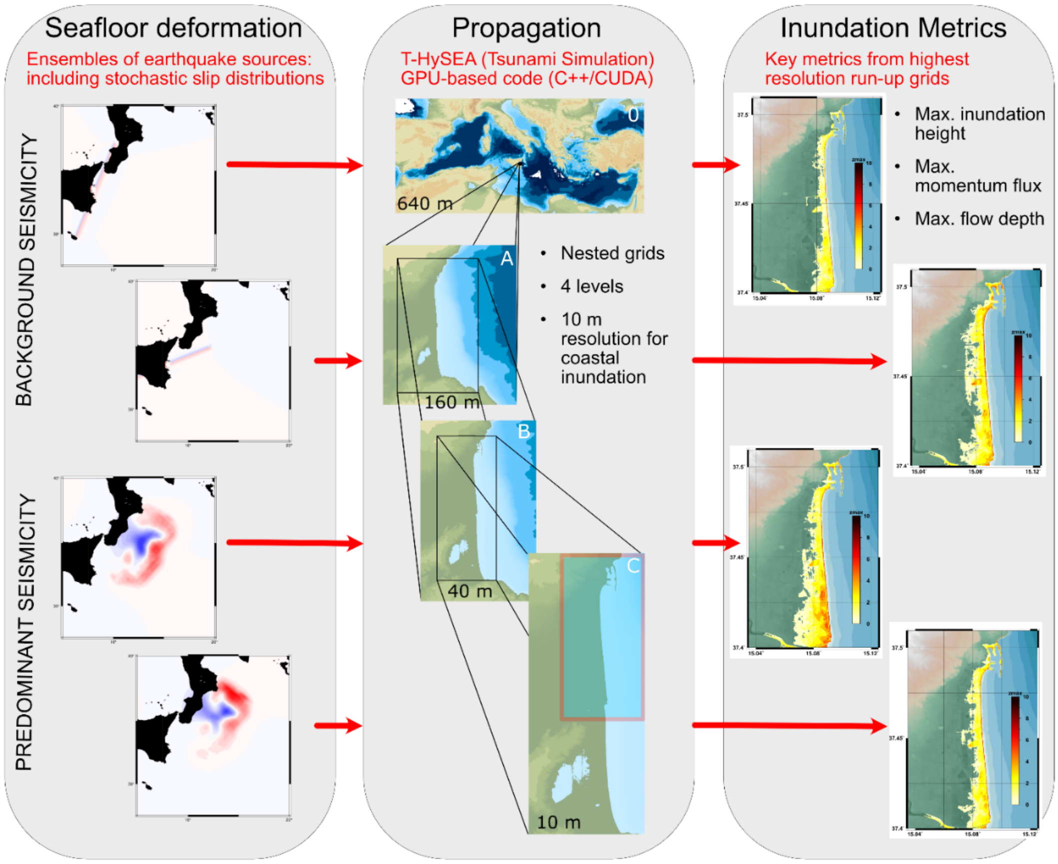


Figure 7.2. Overview of simulation process, from source models and seafloor displacements (left) for BS and PS sources, the nested grid procedure, and tsunami inundation for different scenarios.

**Availability of software:** The process of distributing the software has started and we expect that it will be distributed before the end of the ChEESE project. A software service page has been developed to make the HPC workflow generally available under a Creative Commons license. Furthermore, the access will be facilitated through a dedicated portal built by the candidate Thematic Core Service Tsunami (cTCS-Tsu) in the framework of EPOS-ERIC infrastructure, that enables long term access sustainability and ensures standardization towards other hazard components. The services of the cTCS-Tsu have just started and will soon be available through the community portal https://tsunamidata.org/ (presently under construction) created and maintained by INGV. This is a gateway to all the service portals provided and maintained by the cTCS-Tsu partners across Europe. Selected services (e.g. NEAMTHM18 which is the basis for PD7) are also being made interoperable with those from the other EPOS TCSs via integration in the ICS (https://www.ics-c.epos-eu.org/).

Description of the service, codes will be available here at a later stage:

https://www.ngi.no/eng/Services/Technical-expertise/Tsunamis/HPC-enabled-local-probabilistic-tsunami-hazard-assessment-workflow.

This portal will be accessible through the cTCS-Tsu and through the ChEESE website where a repository of workflow metadata will be created.

**Use case #1. Local PTHA for Catania and Siracusa**

***Summary of technological achievements***

Weak scalability of the flagship code *Tsunami-HySEA* is imperative for efficient simulations on the Tier-0 system. Regarding the weak scalability, this was found to be optimal as could be expected from the minimal amount of I/O performed during the simulations, a task that is performed mostly at the end of the process. One recent implementation of *Tsunami-HySEA* (version 3.8.1MC) has been produced specifically for performing a large number *N* of synchronous simulations as a single job using *N* GPUs. If the earthquakes selected in a single job are close to each other (similar source locations and parameters), the computational requirements for the different simulations are similar and the loss in scalability is minimal (around 1%). If the scenarios are selected randomly, but attending to a predefined order, the loss is between 2-3%. The largest loss measured for the scenarios simulated during the project for Sicily (i.e. the worst-case situation) is below 8%. This allowed us, with no extra effort, to remain below a 2-3% loss in scalability when using 64 GPUs. Table 7.1 presents the figures for weak scalability up to 64 simulations in 64 GPUs (data obtained in Nvidia V100 Graphic Cards in the CTE-POWER cluster at BSC for East Sicily test cases). For the test cases of Catania and Siracusa, **the typical job was run on 128 GPUs**, on the machine **Marconi-100**. The system was also tested at the full capacity allowed by the PRACE project TsuHazAP, **up to 1024 concurrent GPUs**. The scalability was not tested for these higher numbers of GPUs, but is believed to be slightly smaller than for 64 concurrent GPUs.

| # exp / # GPUs used (N) | 1 | 2 | 4 | 8 | 16 | 32 | 64 |
|---|---|---|---|---|---|---|---|
| Time 1 exp / Time N exp | 1 | 0.9994 | 0.9992 | 0.9984 | 0.9968 | 0.9947 | 0.9896 |

Table 7.1: Weak scalability of Tsunami-HySEA

***Scientific achievements***

The following scientific achievements are completed:

- Implementation and testing of the local PTHA workflow in an operational environment. In addition to Catania and Siracusa, a PTHA was performed for Colonia Saint Jordi, Heraklion, Messina, Thessaloniki, and a location in South East Iberia. A total of 1,842,296 high resolution tsunami simulations were performed for the PTHA for the sites listed above. The conclusions from the other test sites are similar as for Catania and Siracusa.
- Carried out a first local PTHA for Catania and Siracusa (Gibbons et al., 2021).
- Addressed uncertainty in the definition of inundation zones for evacuation under a tsunami warning and for long-term coastal planning (Tonini et al., 2021).

- Carried out a large range of sensitivity tests related to PTHA for the sites listed above and, in addition, for Sines in Portugal. Many results are yet to be analysed, although some are available in a soon-to-be-published paper (Gibbons et al., in final revision). The sensitivity analysis gives additional insight into the PTHA results and their uncertainty and helps to design PTHA better in the future.

***Scientific Products***

***Papers***

- Gibbons, S. J., Lorito, S., Macías, J., Løvholt, F., Selva, J., Volpe, M., ... Vöge, M. (2020). Probabilistic tsunami hazard analysis: high performance computing for massive scale inundation simulations. Frontiers in Earth Science, 8:591549, doi: 10.3389/feart.2020.591549.
- Gibbons, S.J., Lorito, S., de la Asunción, M., Volpe, M., Selva, J., Macías, J., Sánchez Linares, C., Brizuela, B., Vöge, M., Tonini, R., Lanucara, P., Glimsdal, S., Romano, F., Christian Meyer, J., and Løvholt, F. (in revision), The Sensitivity of Tsunami Impact to Earthquake Source Parameters and Manning Friction in High-Resolution Inundation Simulations, Frontiers in Earth Science, in final revision.
- Løvholt, F., Glimsdal, S., and Harbitz, C. B. (2020). On the landslide tsunami uncertainty and hazard. Landslides, 17, 2301-2315. doi: 10.1007/s10346-020-01429-z
- Tonini, R., Di Manna, P., Lorito, S., Selva, J., Volpe, M., Romano, F., Basili, R., Brizuela, B., Castro, M.J., de la Asunción, M., Di Bucci, D., Dolce, M., Garcia, A., Gibbons, S.J., Glimsdal, S, González-Vida, J.M., Løvholt, F., Macías, J., Piatanesi, A., Pizzimenti, L., Sánchez-Linares, C., Vittori, E. (2021). Testing Tsunami Inundation Maps for Evacuation Planning in Italy. Frontiers in Earth Science, 9:628061. doi: 10.3389/feart.2021.628061.

***Conferences, seminars***

- Gibbons, S.J., Castro, M.J., Glimsdal, S., Harbitz, C.B., Lorenzino, M.C., Lorito, S., Løvholt, F., Nazaria, M., Romano, F., Macías, J., Selva, J., Tonini, R., González-Vida, J.M., Volpe, M., and Vöge, M.: Probabilistic Tsunami Hazard Analysis: High Performance Computing for Massive Scale Monte Carlo type Inundation Simulations, EGU General Assembly 2020, Online, 4–8 May 2020, EGU2020-8041, https://doi.org/10.5194/egusphere-egu2020-8041, 2020
- Gibbons, S.J., Lorito, S., de la Asunción, M., Volpe, M., Selva, J., Macías, J., Sánchez-Linares, C., Vöge, M., Tonini, R., Lanucara, P., Glimsdal, S., Meyer, J. C., Romano, F., and Løvholt, F.: The Sensitivity of Tsunami run-up to Earthquake Source Parameters and Manning Friction Coefficient in High-Resolution Inundation Simulations, EGU General Assembly 2021, online, 19–30 Apr 2021, EGU21-14159, https://doi.org/10.5194/egusphere-egu21-14159, 2021.
- Løvholt, F., Gibbons, S., Lorito, S., Volpe, M., Selva, J., and Macías, J. (2021). High Performance Computing for Probabilistic Tsunami Hazard Analysis,

Presentation in the PASC Minisymposium, Advances in Computational Geosciences, Part III, https://pasc21.pasc-conference.org/program/schedule/presentation/?id=msa153&sess=sess125

**Validation.**

| Functional requirement | Target (from D5.1) | Achieved | Validated (YES/NO) |
|---|---|---|---|
| Single scenario 8 hours wave propagation and inundation solved within 1 GPU hour | O($10^5$-$10^6$) scenarios may be foreseen, with O($10^3$) GPU cores this can be accommodated within a few days to a few weeks' time. | Simulations for the four levels of nested-grid domains for Catania and Siracusa calculated 4 hours of wave propagation in 23 minutes (faster than real time).<br><br>Simulations for the single nested grid sites: Messina, Heraklion, Thessaloniki, Saint Jordi, South East Iberia, and Sines took typically between 10 and 15 minutes for 4-hours of wave propagation (faster than real time), with the computation time crucially depending on the dimensions of the grid with the highest resolution.<br><br>A PTHA for most sites (including source refinement) typically required up to $3*10^5$ scenarios and a total of 1,842,296 simulations were performed for the local PTHA for all sites. | YES |
| Required spatial resolution of the order 10 m for inner grid for single runs | Tests will be run at 10 m inner grid resolution | This was performed for all sites. | YES |
| Required spatial resolution of maximum 30" (~900 m) global grid for single runs | Typical range would cover a 10-25 km coastal stretch | The spatial extent of the highest resolution grids varied somewhat for the different sites, but all covered at least a city-scale region. For all sites other than Catania and Siracusa, the maximum height and maximum momentum flux were written out both for the finest (10m) resolution grid and the much larger 160 m grid, | YES |

| | | | |
|---|---|---|---|
| | | typically covering several tens of km. | |
| Required spatial resolution of maximum 30" (~900 m) global grid for single runs | Tests will be run at 640 m global grid resolution covering the Mediterranean Sea | All simulations were performed with a coarsest scale grid of 640m. | YES |
| Telescopic grid refinement ratio 4 | Respective resolutions, 640 m, 160 m, 40 m, 10 m. | This resolution was employed for all simulations performed. | YES |
| Four levels of topo-bathymetric maps in netCDF format for telescopic grids required as input | Inundation computed in the local higher resolution grids | The 4-level nested grids were employed for all calculations with all spatial data stored in netCDF format. | YES |
| Hazard level reproduced to 95% accuracy of original assessment | For a pre-exascale test we expect this to require O($10^5$-$10^6$) scenarios | A 98% accuracy of the hazard level reproduction was used. | YES |
| Must be based on TSUMAPS-NEAM assessment | All sources and probabilities originate from this assessment | TSUMAPS-NEAM formed the basis for all considerations of the source definitions and probabilities. | YES |
| Provide the hazard maps at a spatial resolution of 10 m for full set of exceedance probabilities for all wet areas | Different hazard metrics, flow depth, tsunami height, momentum fluxes. | Tsunami height (from which flow depth can also be calculated) and momentum flux were written out on spatial grids for all calculations. In addition, time-series at selected locations were written out with tsunami heights and velocities. | YES |

Table 7.2. Validation criteria for PD7.

**Involvement of end-users**

A prototypal version of the PD7 workflow was used already in a study (Tonini et al., 2021) performed by the scientists of the Centro Allerta Tsunami (CAT) of INGV, in collaboration with those of NGI, UMA, ISPRA (the Italian Institute for Environmental Protection and Research) and with representative of DPC (the Italian Department of Civil Protection). The study was jointly supported by the ChEESE project and by DPC itself, in the framework of the national activities aimed for tsunami risk management.

The PD7 workflow allowed us to test the methodology that the SiAM (the national tsunami alerting system composed by INGV, ISPRA, and DPC) has used to define the inundation maps for evacuation and long-term coastal planning. The study was conducted to understand limits and advantages of the current methodology, and for addressing to what extent it can be updated through local Seismic Probabilistic Tsunami Hazard Analysis (SPTHA) studies.

Based on a given level of acceptable risk, Italian authorities in charge of this task recommended considering, as design hazard intensity, the average return period of 2500 years and the 84th percentile of the NEAMTHM18 regional hazard model uncertainty. Safety factors based on analysis of run-up variability and an empirical coastal dissipation law on a digital terrain model (DTM) were applied to convert the regional hazard into the design run-up and the corresponding evacuation maps with a GIS-based approach. Since the regional hazard cannot fully capture the local-scale variability, this simplified and conservative approach is considered a viable and feasible practice to inform local coastal risk management in the absence of high-resolution hazard models.

In our study, two locations on the coast of eastern Sicily were considered, and the local hazard was addressed with the PD7 workflow with the same seismic model as the regional one, but using a higher-resolution DTM and massive numerical inundation calculations with the GPU-based *Tsunami-HySEA* nonlinear shallow water code. This study shows that the GIS-based inundation maps used for planning deal conservatively with potential hazard underestimation at the local scale, stemming from typically unmodeled uncertainties in the numerical source and tsunami evolution models. The GIS-based maps used for planning fall within the estimated “error-bar” due to such uncertainties. The analysis also demonstrates the need to develop local assessments to serve very specific risk mitigation actions to reduce the uncertainty.

More generally, the presented case-studies highlight the importance to explore ways of dealing with uncertainty hidden within the high-resolution numerical inundation models, e.g., related to the crude parameterization of the bottom friction, or the inaccuracy of the DTM. This last part is now being further addressed by the aforementioned new sensitivity study by Gibbons et al. (in revision) with the ChEESE PD7 workflow.

We believe that these practical applications, conducted in collaboration with the actual stakeholders and the decision-makers in charge for tsunami risk management demonstrate the operational relevance and the level of maturity of this workflow.

**Impact**

*PD7 has enabled a major scientific advancement: enabling a PTHA workflow that can work efficiently on GPU-based Tier-0 machines to simulate numbers of scenarios 1-3 orders of magnitude larger than previous investigations for local PTHA.*

Previous PTHA methods had several limitations. Large scale PTHA applications (for example the region-wide TSUMAPS-NEAM assessment) were limited to offshore hazard. Related hazard products (e.g. tsunami flow depth, inundation height and momentum flux) were not resolved at a sufficient resolution for local hazard analysis. In the cases where local hazard has been analyzed with probabilistic methods, it has often been limited to a small number of scenarios, not able to resolve the large uncertainty relevant for PTHA. Moreover, a workflow enabling massive scale HPC for local PTHA has been lacking. In PD7, effective implementation of a local scale PTHA workflow operable on Tier-0 HPC systems has closed this gap. Altogether, this provides several key impacts, both scientifically and operationally:

- Scientifically, by providing a unique benchmark case with an unprecedented number of scenarios spanning a large uncertainty space
- Scientifically, by providing a wide basis for uncertainty quantification
- Operationally, providing a new workflow for resolving PTHA uncertainty to a much higher level than previously available.
- Operationally, providing a potential service for stakeholders in the NEAM region, where users can plug in their own high resolution grids to produce the next generation of hazard maps with close to full uncertainty quantification.

*PD7 is already being applied in the framework of tsunami risk assessment, reduction and management in Italy by the Institutions and the authorities in charge.*

As described in the previous Section, the coastal planning in Italy is already receiving more scientific information which in turn will serve as a guidance for future updates. Since CAT-INGV is also a Tsunami Service Provider in the NEAMTWS (the Tsunami Early Warning and Mitigation System in the North-eastern Atlantic, the Mediterranean and connected seas), the results obtained thanks to PD7 already attracted the attention of several participants in the NEAMTWS, and its application is being considered by other countries.

*PD7 will be distributed as a service by the new TCS Tsunami in EPOS.*

This achievement also demonstrates the impact of ChEESE, since this is one of the first HPC-based services for hazard assessment to enter the EPOS landscape. In turn, this is favouring the finalization of the design and the implementation of the planned EPOS distributed services.

## PD8. Probabilistic Tsunami Forecast (PTF) for Early Warning and Rapid Post Event Assessment

| PD8 | Probabilistic Tsunami Forecast (PTF) for Early Warning and Rapid Post Event Assessment |
|---|---|
| Leader | Stefano Lorito (INGV) |
| Participants | • INGV (Manuela Volpe, Jacopo Selva, Fabrizio Romano, Roberto Tonini, Fabrizio Bernardi, Maria Concetta Lorenzino, Angela Stallone, Stefano Lorito)<br>• UMA (Carlos Sanches-Linares, Jorge Macias, Marc de la Asuncion, Jose Manuel Gonzalez Vida, Manuel J. Castro)<br>• NGI (Steven Gibbons, Finn Løvholt, Malte Vöge, Sylfest Glimsdal)<br>• CINECA (Silvia Giuliani, Isabella Baccarelli, Piero Lanucara)<br>• HLRS (Alexey Cheptsov)<br>• UniNA (Antonio Scala)<br>• GFZ (Andrey Babeyko) |
| Workflow | Suite of bash, python/Matlab, C, GMT codes for workflow execution, pre- and post-processing (https://gitlab.rm.ingv.it/; on the intranet and users need authorization);<br>HLRS WMS-light workflow |
| Engine | Tsunami-HySEA |
| TRL initial | 3 |
| TRL target | 6-8 |
| TRL achieved | 7-8* |

**(*)** Expected to be raised to TRL 8 within the project duration, since both engineering for usage within an Urgent Computing experiment and integration with WMS-light workflow manager for orchestration across distributed resources are being achieved.

### HPC Products (available software and workflows)

The tool was extensively described in D4.8. We report here again on the basic functionality of the tool for the sake of self-consistency of this deliverable.

We wish to note upfront that we consider the TRL 7 already achieved, since the PD8 is up and running in the CAT-INGV Tsunami Warning Centre, being in principle usable to create real tsunami alert messages. It will become operational following a multi-stage process including calibration, technical testing, and scientific review phases, under the supervision of the Italian National Civil Protection Department. Once this process will be finalized, we will consider this version of the PD7 workflow at the TRL 8.

PD8 provides a rapid probabilistic forecast of tsunami inundation, following an earthquake offshore or close to the coast, before it actually occurs or before tsunami observations are available. For near-field tsunami early warning (EW) purposes, the large uncertainty about earthquake location and magnitude, as available in the first minutes, are reflected into forecasting uncertainty. For the purpose of supporting rapid post-disaster intervention, for which more time is available, additional source and even tsunami observations in the subsequent phases can be exploited to eventually narrow down the tsunami forecast uncertainty. In D4.8 it was also explained that the probabilistic forecast can be translated into an alert level, which is necessary in the early warning operations to initiate risk mitigation operations. The service also provides an early estimate of the earthquake parameters with their uncertainty when they are not yet available, as a by-product. The schematic of the PD8 workflow is illustrated in Figure 8.1.

There are two versions of the PTF tool, reported below as two different use-cases. Actually, within each of the use-cases we include several applications, since they have been both run for various earthquakes and tsunamis in the Mediterranean Sea and in the Pacific Ocean.

The first version is based on pre-calculated numerical simulations of tsunami scenarios. It works in near-real time in the premises of the CAT-INGV Tsunami Warning Centre (https://www.ingv.it/cat/en/), which is a NEAMTWS Tsunami Service Provider (http://www.ioc-tsunami.org/). It deals with potentially tsunamigenic earthquakes anywhere in the entire Mediterranean Sea (offshore and inland close to the coast). Order of 10-100 k pre-calculated scenarios were run on a 30 arc-sec grid for 8 hours of simulation. The number of scenarios in the ensemble depends on the magnitude of the earthquake and on the desired level of accuracy, that is, roughly speaking, how many standard deviations are explored around the expected values of the seismic parameters. In this configuration, PTF output provides exceedance probabilities for tsunami heights just off the coastline for almost equally spaced points of interest every 20 km along the coasts of the Mediterranean Sea.

The second version is based on simulation ensembles to be run from scratch on large enough HPC clusters in urgent computing mode. Using the Mw 7.0 Samos earthquake as an input, this version of the workflow is being tested with several large scale runs in urgent computing mode on Marconi100 at CINECA. A reservation of 800 nodes each equipped with 4 V100 GPUs has been used for testing the workflow, considering two standard deviations for the source uncertainty, which leads to almost 40 000 scenario simulations. From the first test several bottlenecks have been identified, and a new test at the same scale was successfully carried out. Consequently new scripts are now being integrated into the prototype WMS-light-based workflow developed in collaboration with HLRS. Once this version will be finalized and tested, we will consider the TRL 8 achieved for this version of the PD8.

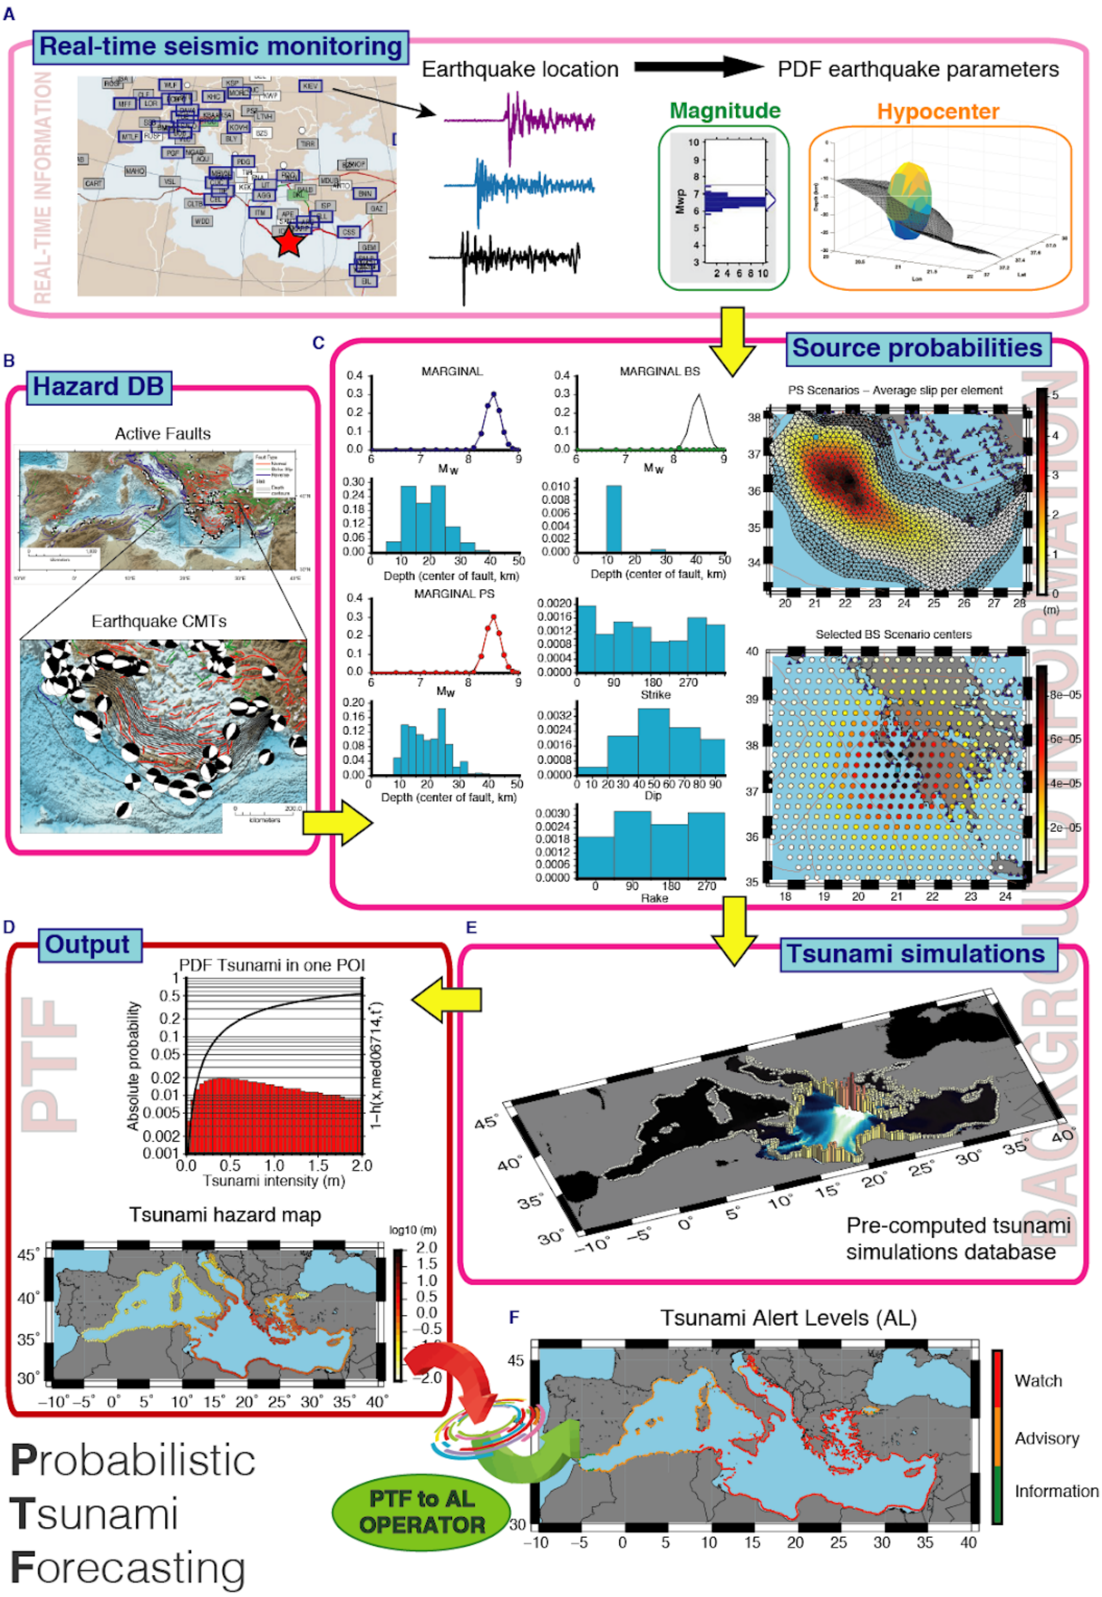


Figure 8.1 Schematic illustration of the PD8 PTF workflow.

Both versions of PD8 progressed from a set of scripts previously run manually to a now fully automated workflow. The version based on pre-computed scenarios can be activated by a Rabbit-MQ message produced by CAT-INGV and can produce automatic alert messages. The urgent computing version needs to be triggered manually but then runs from the pre-processing, through the HPC part, to the post-processing and basic visualisation of the results.

A stand-alone Matlab-based version of the PD8 is available on a github repository (https://github.com/INGV/matPTF) as accompanying material to the scientific paper presenting the methodology (Selva et al., 2021; https://www.nature.com/articles/s41467-021-25815-w), which fully acknowledges the ChEESE contribution.

Furthermore, the access will be facilitated through a dedicated portal built by the candidate Thematic Core Service Tsunami (cTCS-Tsu) in the framework of EPOS-ERIC infrastructure, that enables long term access sustainability and ensures standardization towards other hazard components. The services of the cTCS-Tsu have just started and will soon be available through the community portal https://tsunamidata.org/ (presently under construction) created and maintained by INGV. This is a gateway to all the service portals provided and maintained by the cTCS-Tsu partners across Europe. Selected services (e.g. NEAMTHM18 which is the basis for the PD8 version for the Mediterranean Sea) are also being made interoperable with those from the other EPOS TCSs via integration in the ICS (https://www.ics-c.epos-eu.org/).

**Use case #1. PTF based on pre-calculated scenarios - Early Warning Mode**

***Summary of technological achievements (HPC performance, etc)***

The PD8 version in early warning mode was run, retrospectively, for all events dealt with by CAT-INGV in recent years. For example, it was run for the recent Ierapetra and Samos-Izmir earthquakes and tsunamis in 2020. More recently, it was tested in pure forecasting mode for two strong earthquakes in the Eastern Mediterranean. Using the pre-calculated scenarios, the computing time can be limited to the order of seconds, which is suitable for early warning purposes.

The HPC resources are needed in this case to pre-compute the scenario set. The computing time of a single simulation on 1 GPU (V100) is for example ~300 s for 8 hours of tsunami propagation within the entire Mediterranean at 30 arcsec (~ 900 m) spatial resolution, or ~2 hours for 40 hours of propagation within the entire Pacific Ocean at 1 arcmin resolution; the total time to run the whole ensemble depends on the ensemble size and on the number of GPUs available; in this case the pre-computed scenarios are on the order of tens of thousands.

**Use case #2. PTF based on on-the-fly scenarios - Rapid post-event assessment mode (Urgent Computing)**

***Summary of technological achievements (HPC performance, etc)***

This second version was already tested with: a suite of more than 10 recent events in the Mediterranean Sea, that is all those for which the CAT-INGV issued a tsunami alert message in recent years; the 2003 Zemmouri-Boumerdes Mw 6.8 earthquake and tsunami in the western Mediterranean; the NEAMWave17 Mw 8.5 synthetic scenario; the 2010 Mw 8.8 Maule Chile earthquake and tsunami. Also in these cases, orders of 10-100 k scenarios were run on a 30 arc-sec grid for 8 hours of simulation (and 30 hours for the Maule event in the Pacific). Also in this configuration, PTF output provides exceedance probabilities for tsunami heights in front of the coastline, but for denser almost equally spaced points of interest every 2 km along the coasts. For the Maule event, the output is also retrieved in correspondence of the deep sea DART sensors that measured the tsunami. It is now being tested for an urgent computing exercise organized by the ChEESE project with the Mw 7 2020 Samos earthquake, on a local grid (see for example deliverable 5.8). As noted above, almost 40 000 scenarios are necessary for exploring the source uncertainty up to 2 standard deviations. In this configuration, the test took less than 10 minutes on 800 Marconi100 nodes.

The flagship code for the individual tsunami numerical simulations, Tsunami-HySEA, has been tested on many different supercomputers, such as CTE-POWER (BSC), DAVIDE and Marconi100 (CINECA), Piz Daint (CSCS), HPC4 and HPC5 (ENI), in different frameworks, among which is worth noting the Project TSU-CAST - TSUnami ForeCASTing in the PRACE Call 20. https://prace-ri.eu/hpc-access/project-access/project-access-awarded-projects/projects-awarded-under-prace-project-access-call-20/. The reported weak scalability reported for PD7 above is also relevant here.

Some computational details for use-case #2 are reported in the table below. Hence, they refer to the post-event assessment only, without inundation modelling, on a rather limited spatial domain, and for a limited duration of the simulation: to consider Early Warning mode and/or inundation modelling and/or a very large domain (e.g. the Pacific Ocean) would possibly require exa-scale capacity.

| | **Number of GPUs:** | **Memory (GB):** | **Storage (GB) both temporary and permanent** | **#files written both temporary and permanent** | **I/O data traffic per hour during job** |
|---|---|---|---|---|---|
| Minimum*: | **256** | **16 GB/GPU** | **80GB** | **40k** | **26GB** |
| Average*: | **512** | **16 GB/GPU** | **2TB** | **40k** | **600GB** |

| | | | | | |
|---|---|---|---|---|---|
| Maximum*: | **1024-3200**** | **32 GB/GPU** | **4TB-500GB**** | **40k** | **1.3TB-N/A**** |

* Min to Max refers to increasing domain size and simulation duration for an application dealing with Rapid Post event assessment performed in a target time of 3 hours with an ensemble size of 20000 scenarios.

** Using the maximum is recommended by the HPC centers. So far, the 1024 number of contemporaneous GPUs were limited by the availability of resources through PRACE. This number was increased to 800 nodes with 4 GPU each on Marconi100 at CINECA in the last tests. In this case the output size was reduced, further tests are ongoing, and the traffic is still being measured.

***Scientific achievements***

As stated on twitter by Nature Communications, "An article published in @NatureComms introduces Probabilistic tsunami forecasting — an approach to tsunami early warning that quantifies uncertainty, enhances forecast accuracy and enables rational decision making.".

We consider this an important scientific achievement, which paved the way for the operation use of PD8.

The PD8 was tested against several real events, both in the version based on pre-calculated scenarios and the version based on on-the-fly simulations in urgent computing mode; both at the Mediterranean and at the Pacific Ocean scales. The results of some of these tests are reported in the Selva et al. (2021) Nature Communications paper.

Some of these tests have been conducted in the framework of the PRACE Project TSU-CAST, awarded to PD8 on the CINECA Marconi100 supercomputer. The PRACE Project allowed us to perform further tests, whose results are now being extensively analyzed and will be the basis for future publications.

Seismic and tsunami forecasting are being used for joint PTF calibration against earthquake and tsunami observations, both in the Mediterranean Sea and in the Pacific Ocean, and for a wide range of earthquake magnitudes and tsunami intensities. This is a necessary condition for allowing its transition to the status of an operational tool for tsunami early warning.

***Scientific Products***

- Lovholt, F. et al. Urgent Tsunami Computing. in 2019 IEEE/ACM HPC for Urgent Decision Making (UrgentHPC) 45–50 (IEEE, 2019). doi:10.1109/UrgentHPC49580.2019.00011.
- Selva, J. et al. Probabilistic Tsunami Forecasting (PTF) for Tsunami Early Warning operations. Geophysical Research Abstracts Vol. 21, EGU2019-17775, 2019. EGU General Assembly 2019, solicited.
- Selva, J., Lorito, S., Volpe, M. et al. Probabilistic tsunami forecasting for early warning. Nat Commun 12, 5677 (2021). https://doi.org/10.1038/s41467-021-25815-w.


**Validation.**

| **Functional requirement** | **Target (from D5.1)** | **Achieved** | **Validated (YES/NO)** |
|---|---|---|---|
| Interface with real-time seismic/tsunami monitoring system | Implement the PD8 in the CAT-INGV system | YES | YES |
| Access to TSUMAPS-NEAM database | Expand the storage of CAT-INGV system; transfer the simulation database | YES | YES |
| Global regionalization, subduction models and probability distribution of earthquake mechanisms | Evaluate the prior probabilities for earthquakes in the Pacific | YES | ONGOING |
| Global amplification factors | Implement amplification factors everywhere to evaluate inundation | NO - a more basic approach based on Green's law temporarily adopted | NO |
| Global tide-gauge/DART locations, and data stream | Implement the positions of the instruments in the simulation setup in the Pacific Ocean | YES | YES |
| Topo-Bathymetric data | Implement the DEM in the simulation setup in the Pacific Ocean | YES | YES |

**Involvement of end-users**

Stakeholders already involved in the development phase are: the CAT-INGV NEAMTWS Tsunami Service Provider (http://www.ingv.it/cat/en/), the Italian National Civil Protection Department (http://www.protezionecivile.gov.it/home), the

ARISTOTLE-eENHSP multi-hazard scientific partnership (http://aristotle.ingv.it/tiki-index.php).

CAT-INGV is an operational tsunami warning that defines, as such, the time requirements to issue an alert, which are a few minutes. Since this time, potentially used for the population to escape a tsunami, is already partially consumed by the acquisition of long enough seismic time-series and their processing, plus the necessary checks by the personnel on shift regarding the credibility of the automatic seismic solution and the correctness of the alert messages prepared by the system, the time to elaborate the forecast with the PTF needs to be reduced to a few seconds. When the calibration and revision process will be finalized the PTF will become operational in a tsunami warning centre. This would be the first time that an official warning centre adopts this new methodology.

On the other hand, the ARISTOTLE requirements are less strict time-wise. ARISTOTLE needs to report within one-three hours after the first registration of the event regarding the expected impact. Hence, higher resolution and better accuracy might be gained in case sufficient computational resources are available, compared to the early warning case.

Potential future stakeholders are the members of the tsunami scientific community, as the PD8 can be used to analyse retrospectively past tsunami events, for example.

It can also be expected that other tsunami warning centres will adopt the PTF or similar methods in the future.

**Impact**

*Quantification of uncertainties in real-time tsunami forecasting.*

We present a novel approach dealing with uncertainty in real-time tsunami forecasting, coined Probabilistic Tsunami Forecasting (PTF), for use in tsunami early-warning, but with application also for rapid post-event assessment. Existing practices are either deterministic or only consider the uncertainty implicitly and in a rudimentary way.

*Separation of the roles between scientists and decision makers.*

PTF allows a clear distinction between the forecast of the impact of the ongoing tsunami and the decision to be taken to reduce the associated risk. This can be done in a transparent manner based on criteria established in advance by the decision makers in combination with the PTF.

*Uncertainty quantification is complementary to uncertainty reduction.*

The PTF is a rigorous approach to tsunami forecasting with uncertainty; this is synergistic to the technology-driven uncertainty reduction effort through enhanced real-time tsunami monitoring capability (GNSS, DART buoys, SMART cables). These two elements have been already emphasized by the United Nation Decade of Ocean Science for Sustainable Development (2021-2030, https://www.oceandecade.org/).

*Exploitation of capacity offered by pre-exascale supercomputers.*

Greater computational power permits

1) up-scaling for a deeper uncertainty characterization, as more scenarios could be added to better represent the natural variability, or to characterize the uncertainty associated to numerical modelling;
2) up-scaling to include "more physics" in the simulations. In fact, the simulation schemes are often simplified to reduce the computational cost (e.g. as in PD4);
3) up-scaling to achieve tsunami impact characterization at higher resolution (e.g. as in PD7), on larger geographical domains and longer duration of the simulations. For example, the inundation which is characterized by relatively short wavelengths is not presently simulated with the resolution permitted by currently available resources, and inundation features must be extrapolated from simulations of offshore waves;
4) faster calculations, with any of the above kept fixed.

*Down-scaling for early-warning*

Down-scaling of the computational cost would be useful for developing a light version for early warning. This is being achieved via HPC-based full uncertainty characterization and subsequent sampling/parameterization of the ensemble; this would allow usage also in the absence of huge computational resources in day-by-day warning operations.

*PD8 will be distributed as a service by the new TCS Tsunami in EPOS.*

This achievement also demonstrates the impact of ChEESE, since this is one of the first HPC-based services for hazard assessment to enter the EPOS landscape. In turn, this is favouring the finalization of the design and the implementation of the planned EPOS distributed services.

## PD9. Seismic Tomography

| PD9 | Title |
|---|---|
| Leader | Vadim Monteiller (CNRS) |
| Participants | Amandine Sergeant, Masaru Nagaso |
| Workflow | Suite of bash, python, C++, codes for workflow execution, pre- and post-processing; graphical interface. |
| Engine | SPECFEM3D, SALVUS |
| TRL initial | 4 |
| TRL target | 6 |
| TRL achieved | 6 |

**HPC Products (available software and workflows)**

The purpose of PD9 is subsurface imaging using seismic wave propagation. Seismic waves propagate underground, interact with geological 3D structures and are then recorded at the surface by sensors. The principle of PD9 is to use an inverse problem to adjust the recorded data and the synthetics computed by modeling, using an optimization process to fit the measurements. This is an iterative optimisation process which requires three simulations per seismic source. Thus the total number of simulations required is three times the number of sources (50-300) multiplied by the number of iterations (10-50) in each frequency band (2-10). This leads to several thousands or tens of thousands simulations to obtain the final model. To produce better quality images, PD9 relies on **large amounts of data**, **accurate wave physics modelling** and **high frequency** seismic modelling.

The **PD9 workflow** consists of three parts: 1) a **pre-processing** which consists in selecting the data, defining the mesh and the initial model; 2) the **calculation of the inversion** itself; and 3) a **post-processing** to visualize the results, the models and the modeled synthetics.

The first part of the workflow is composed of two parts, the first one consists in downloading the data and changing their format, the second part consists in doing a quality control of the data and extracting a database usable for the inversion. Since we use the complete seismogram to do the inversions, contrary to classical methods that use only a small part of the seismic attributes, we have to be very careful on the data selection to avoid introducing artifacts in the final images. For this purpose we have developed graphical tools to help the user in this crucial step. A tool for seismology in Python and a tool for geophysical exploration in C++. The quality control of seismograms is based on metrics that depend on the acquisition geometry, the frequency content and the wave propagation regime. The demands for the quality controls are not exactly the same in both cases. Hence we choose to make different tools for the two application domains.

The inversion (second part) is performed by using the open-source **ChEESE Flagship code SPECFEM3D** (https://github.com/geodynamics/specfem3d) on a computing cluster, as this is the part of the workflow that requires the most computing resources. The **SALVUS code** has also been tested on the use-cases. Several iterations of the inversion process are launched, generally between 10 and 50, for each chosen frequency band. The frequency is gradually increased in order to obtain the expected spatial resolution. The mesh must also be adapted to the frequency band, and is generally refined as the frequency increases, requiring increasing computing resources with increasing frequencies.

The third part, the post-processing, consists in verifying the good fit of the data and synthetics in the final model and in visualizing the models for their geological interpretation.

**Use case #1. Ocean Bottom Seismometer in North sea**

***Data***

For this first case, we have data from Equinor, a Norwegian energy company that has made public data from the Volve oil field in the North Sea (https://www.equinor.com/en/what-we-do/digitalisation-in-our-dna/volve-field-data-village-download.html).

Several explorations have taken place on this oil field, and we have chosen to process the data recorded by the Ocean Bottom Seismometers (OBS). A total of 137 instruments were deployed on the seabed and more than 70,000 shots were fired with air guns at 5 meters below the sea level every 25 meters. Thus more than 28 million seismograms were recorded with a sampling frequency of 500Hz. Each seismogram contains 5000 samples, which makes a total data set of over 1TB.

OBS are placed on the seafloor and record the three components (two horizontal and one vertical) of particle velocity (figure 9.1, 9.2). In general, in marine exploration, only pressure is used for imaging subsurface structures. This reduces the number of components to one instead of three. Moreover, this implies that the properties of the earth crust can be described as a fluid. This simplifies the modeling considerably. In our case, **we consider the full complexity of the physics by using a visco-elastic and anisotropic rheology in the subsurface**, coupled with the fluid domain which is the seawater layer.

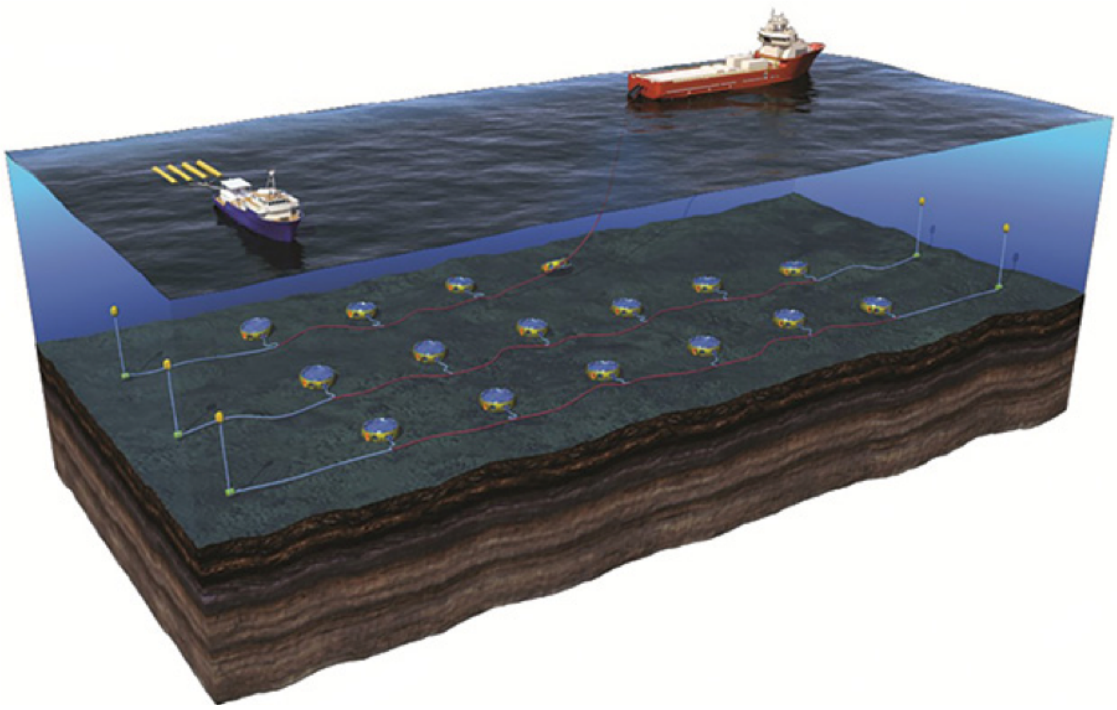

Figure 9.1. Setup for exploration geophysics offshore. The Ocean Bottom Seismometers (OBS) are on the seafloor and record the seismic waves triggered by airguns pulled by boat.

***Quality control***

Data selection must be done before any inversion to keep only the data with an acceptable signal-to-noise ratio in the different frequency bands. To assist the user in the interactive selection of data, we have developed a graphical interface written in C++ and based on the C++ qt5 library. The choice of C++ seemed to us to be the most adapted to be able to treat this large volume of data quickly. The very low latency of the resulting tool allows many fast interactions at the database and graphical display level.

To select the data, we inspected all traces recorded by the OBSs. For each OBS, we have about 6GB of data consisting of three-component seismograms. We proceed component by component, each component has about 2GB of data consisting of over 70,000 seismograms. Shots were triggered every 25 meters by airguns towed by boats using about 150 lines. To visualize the data, we display on the screen each line sorted in the order of the shots, which gives us between 400 and 600 traces displayed per line (figure 9.3). During inspection of the traces, we can delete a trace from the database by a mouse click (figure 9.3). Once the line has been processed, we process the next one, then we process component by component starting with the vertical component followed by the horizontal components. We have encountered several cases where we have to delete a trace. Sometimes the traces present events from other seismic sources which are superimposed on the recorded shot which makes the trace unusable.

Sometimes noise likely originating from an engine affects the seismic trace. It also happens that the traces contain impulses randomly distributed, which we interpret as problems coming from the instrument itself. In all these cases, we must eliminate the traces before the inversion. We also had to eliminate some OBSs because they had a majority of traces that were not suitable for waveform inversion. For the final inversion, we kept only 128 OBS and eliminated about 10% of the traces recorded by these instruments.

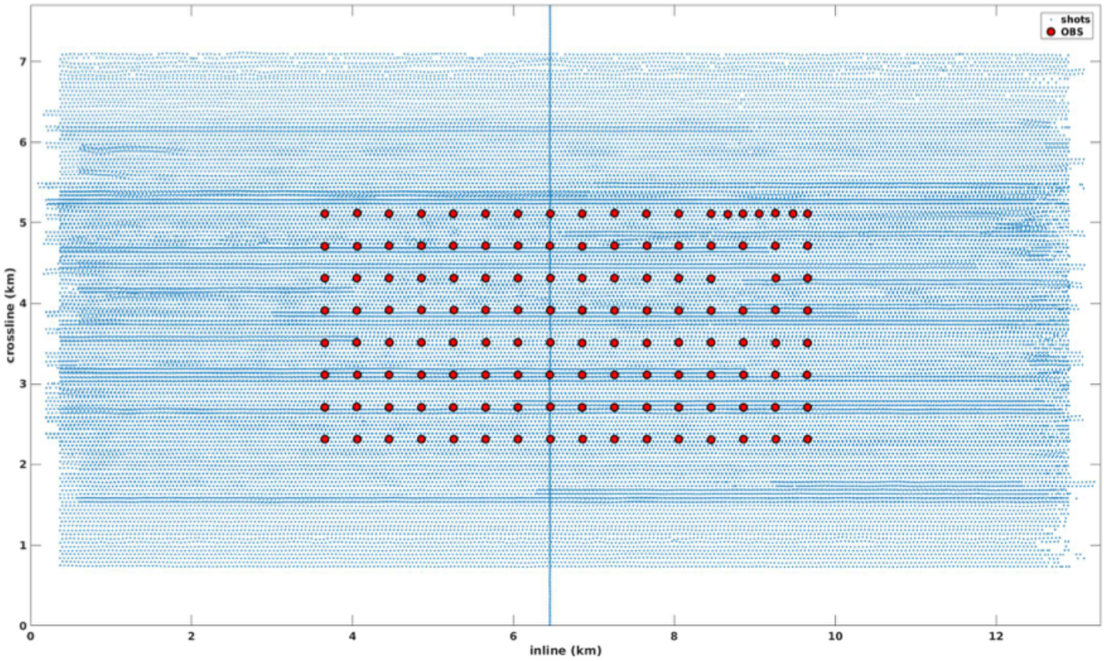


Figure 9.2: Map of acquisition geometry. The blue dots are the position of shots and red circles are the OBS.

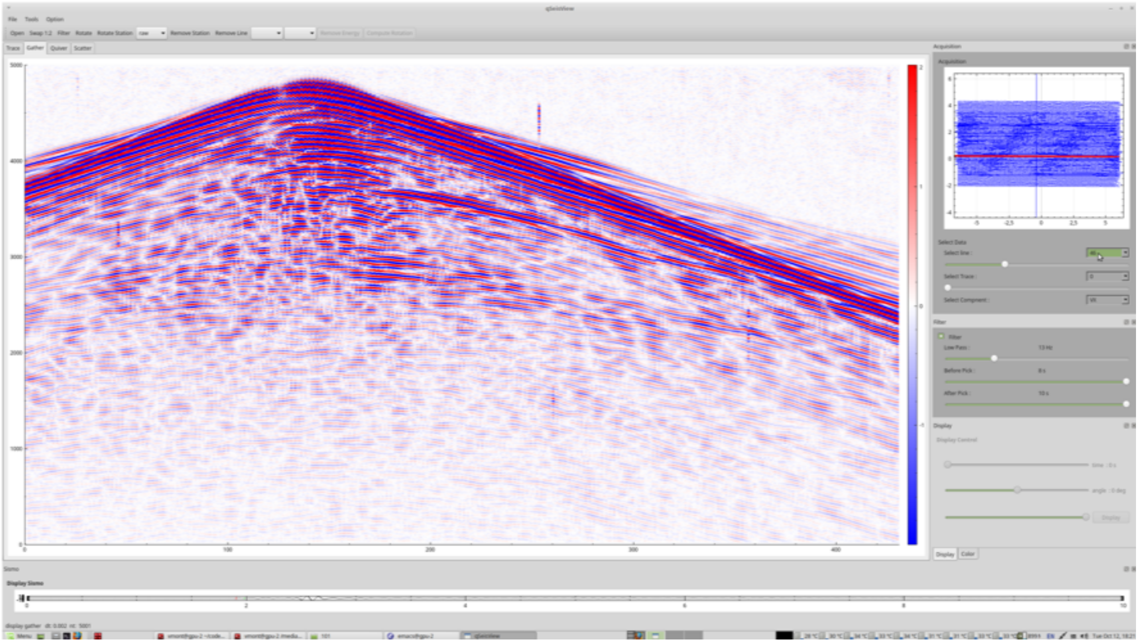

Figure 9.3. Example of records displayed by the graphical interface developed in C++ and qt5. The main display shows ~500 traces recorded by one OBS for the line in red (top right). We clearly see a spurious trace which can be deleted by a mouse selection.

***Mesh and initial model.***

The initial model was computed using a conventional travel time tomography method in exploration geophysics. This allows us to start the waveform inversion in a smooth model that contains the long wave structures. We have six parameters, density, P-wave velocity, S-wave velocity and three Thomsen parameters describing the vertical transverse anisotropy. We also added two quality factors describing the attenuation of P- and S-waves. The characteristics of the initial model allowed us to define a meshing strategy. We used three layers of elements of different sizes (figure 9.4). The first layer contains the sea water and sediment layer which extends 500 metres below sea level. The second layer below the sediment layer contains elements three times larger than the previous ones and extends to 2800 meters depth. The last layer contains elements twice as large and stops at a depth of 5000 metres. The size of the elements is adjusted according to the frequency range used for inversion.

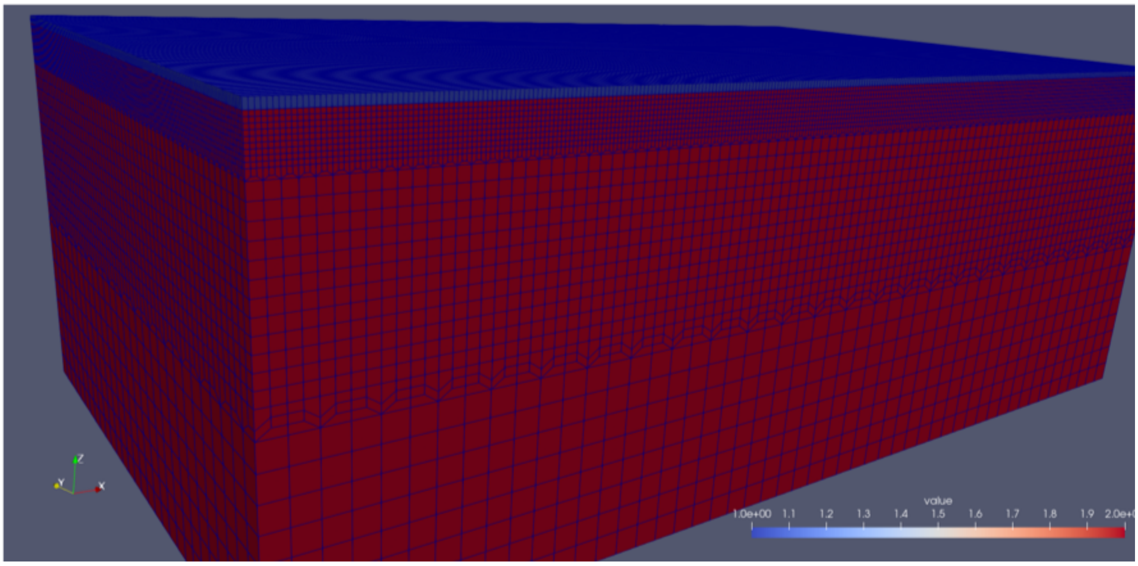


Figure 9.4 Mesh of the domain consists of 3 layers with different element size according to the seismic wave velocity.

***Inversion process: technological achievements***

For this use case, we used the Marconi100 (CINECA) GPU cluster using the allocated computational hours from the PRACE project : 2019215212. The inversion was performed in a hierarchical manner by increasing the maximum frequency. The mesh size and number of time steps were adapted to each frequency band. We performed the inversion in 8 frequency bands. We started the inversion between 0.1 and 4.5 Hz and ended up between 0.1 and 9 Hz (Table 9.1). To perform these inversions, we had to use up to **1024 Nvidia V100 GPUs and a total of 223100 GPU.hours**. The iterations in each frequency band were stopped when local minimum of the misfit function was reached.

| Frequency (Hz) | Number of elements | Number of time steps | Number of iterations | Number of GPU | GPU.hours |
|---|---|---|---|---|---|
| 0.1 - 4.5 | 200000 | 13500 | 52 | 512 | 4600 |
| 0.1 – 5 | 200000 | 13500 | 26 | 512 | 2200 |
| 0.1 – 5.5 | 200000 | 13500 | 16 | 512 | 1300 |
| 0.1 – 6 | 320000 | 21600 | 46 | 640 | 41000 |
| 0.1 – 6.5 | 320000 | 21600 | 25 | 640 | 22000 |
| 0.1 - 7 | 400000 | 21600 | 30 | 512 | 43000 |
| 0.1 - 8 | 730000 | 27000 | 30 | 1024 | 86000 |

| | | | | | |
|---|---|---|---|---|---|
| 0.1 - 9 | 730000 | 27000 | 8 | 1024 | 23000 |

Table 9.1. Frequency band used in hierarchical inversion and computational resources.

***Results***

We considered three elastic parameters for the inversion, density, P-wave velocity and S-wave velocity.Each parameter is updated during the iterations, which allows a better adjustment of the data (Figure 9.5). The variation in physical parameters can be interpreted to understand the geological structures of the subsurface. The images essentially show reflectors that correspond to different geological layers. For example, in figure 9.6, between 100 meters and 2 kilometers deep, we see different layers that could correspond to chalk. Between 2.8 kilometers and 3.2 kilometers deep, we see a strong reflector that corresponds to a solid material that would act as a trap for hydrocarbons. Below, between 3.2 and 3.6 kilometers deep, we see slower velocities that could correspond to the hydrocarbon reservoir. These images must then be provided to teams of specialized geologists to make a more accurate interpretation.

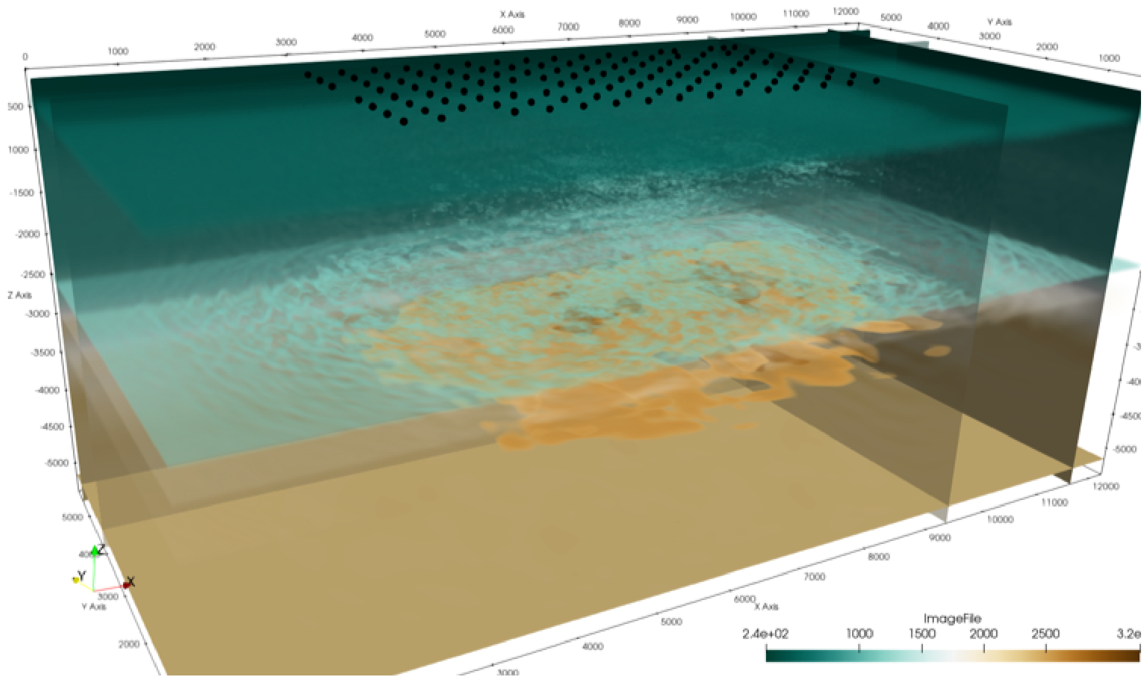


Figure 9.5. 3D geological structures recovered by 3D viscoelastic seismic Full Waveform Inversion. The data are recorded by 128 Ocean Bottom Seismometers in the North Sea (black dots). The color palette represents the different geological layers.

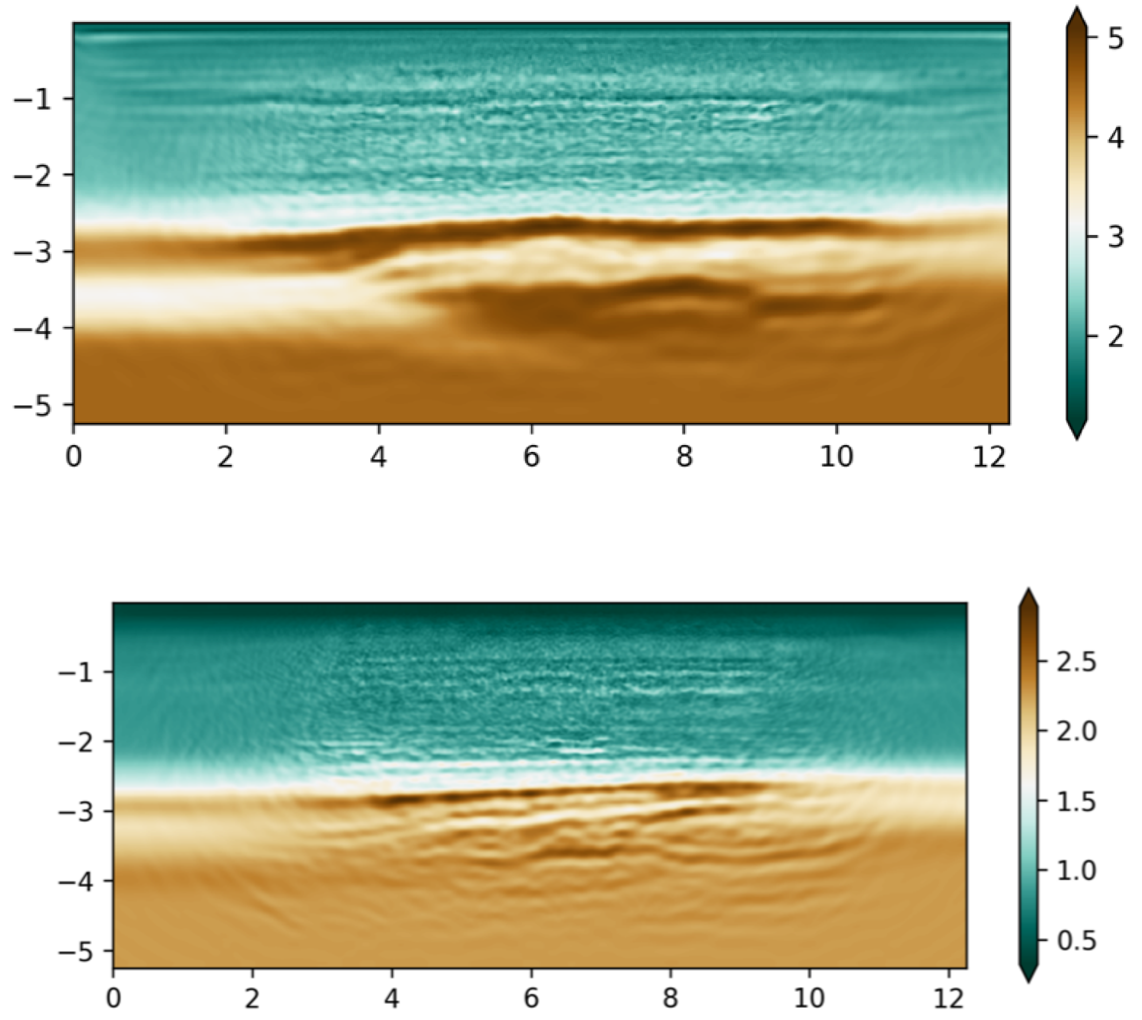


Figure 9.6. Vertical cross-sections of the final model (top P-wave velocity, bottom S-wave velocity) The variation of material properties can be interpreted by geologists to understand the subsurface structures.

**Use case #2 : Iberia-Pyrenees**

***Data fetch***

This use-case is **based on open data** downloadable from different European data centers. We used all available stations in France, Spain and Morocco during the period 2011-2020 (~800 stations at all). Our workflow allows us to make direct requests to the different data centers and to download potential events for seismic tomography. We performed a first selection on magnitude (>5.5) and epicentral distance (>35° and <100°). Another selection is then made by analyzing the focal mechanisms and the azimuthal distribution of hypocenters. Finally, we selected only the events with a signal-to-noise ratio greater than 5. We thus selected 63 earthquakes for our inversion (Figure 9.7).

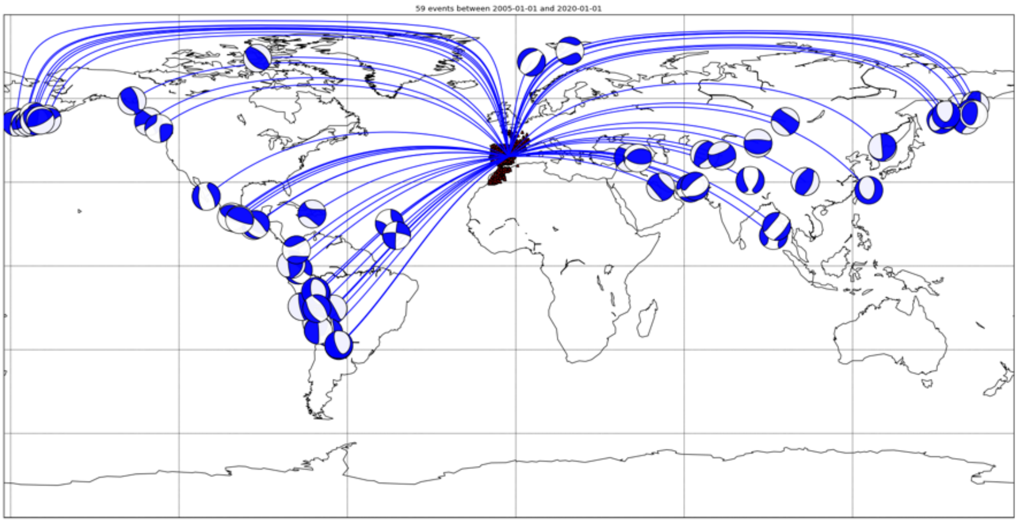


Figure 9.7, Acquisition map showing selected earthquakes and stations.

***Data selection***

Once the earthquakes are selected, we need to inspect the traces to eliminate noisy components or stations with instrumental problems. We have developed a **graphical interface in Python**, which allows us to visualize the traces, the network, to filter the traces and to do different quality checks in order to clean the dataset. We have both automatic criteria to eliminate the traces or we can do it, for some, by a simple mouse click (figure 9.8).

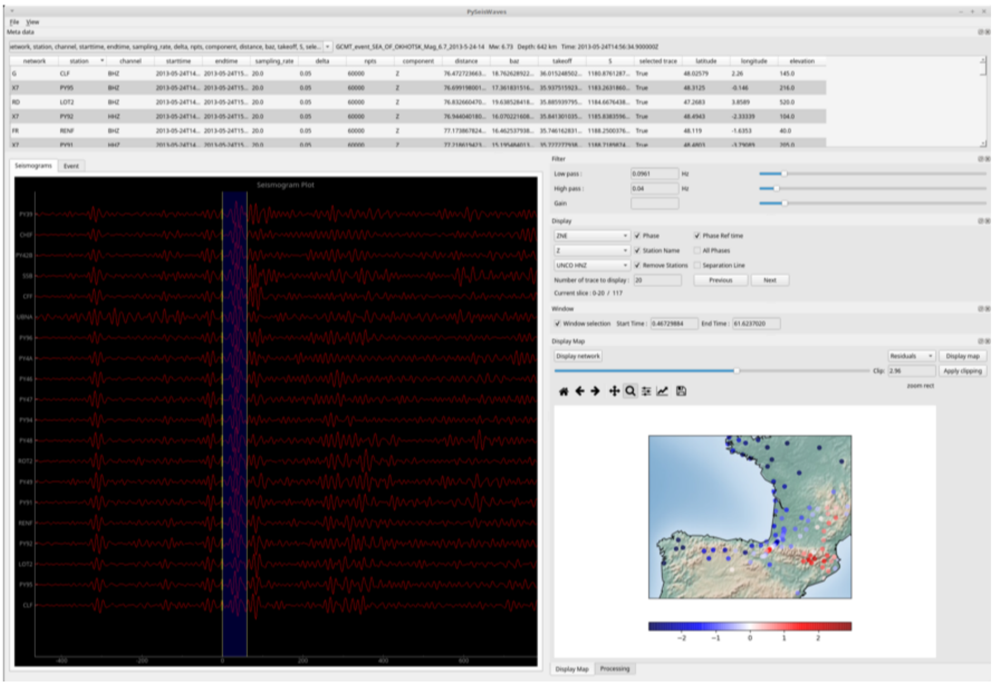

Figure 9.8, Graphical interface allowing to visualize the traces, the different metadata and to make a quality control on the traces. Bad seismic traces are eliminated by using metrics that can be visualized by colors on the map.

***Inversion***

The nature of the data and the scale at which we work (regional) led us to define 2 frequency bands for the inversion, 0.01-0.08 Hz (12.5s - 100s) and 0.01-0.12 Hz (8s-100s). We performed 50 iterations in each frequency band. Table 9.2 shows that it took us 26000 GPU hours to perform the inversion.

| Frequency (Hz) | Number of elements | Number of time steps | Number of iterations | Number of GPU | GPU.hours |
|---|---|---|---|---|---|
| 0.01 – 0.08 | 100000 | 12000 | 50 | 112 | 4000 |
| 0.01 – 0.12 | 337500 | 20000 | 50 | 336 | 22000 |

Table 9.2. Computational resources and simulation settings for the regional scale tomography.

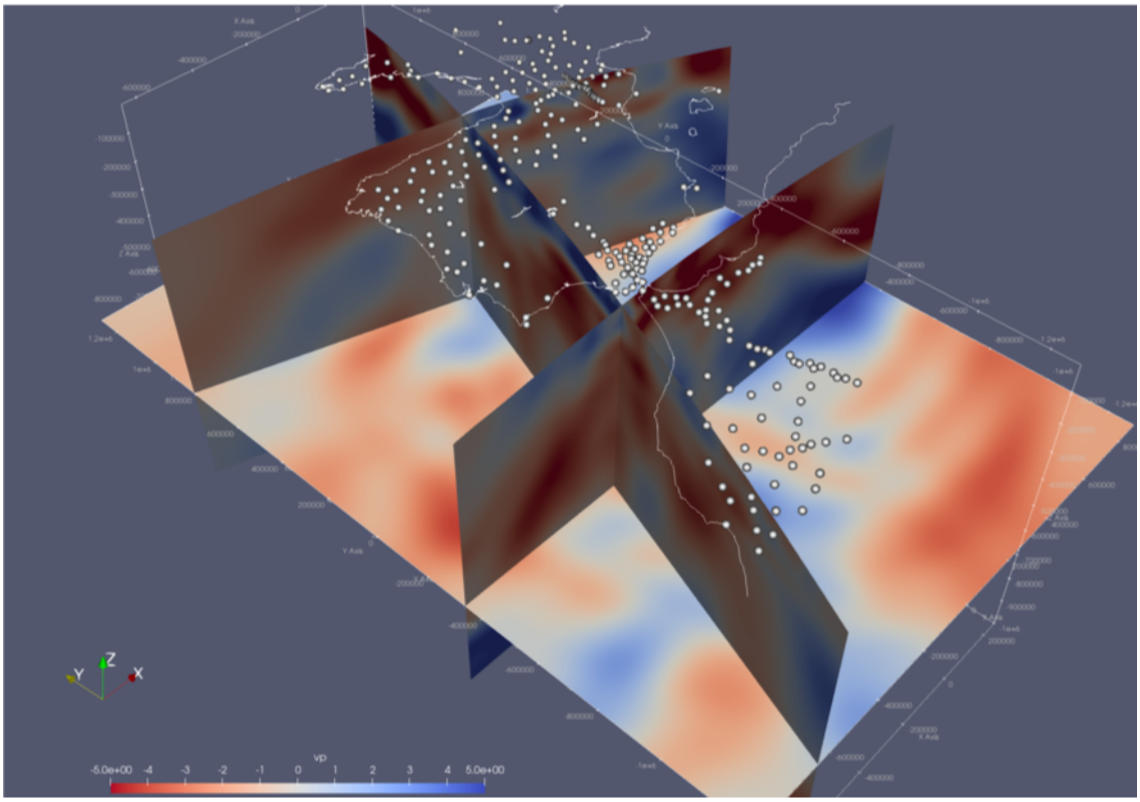

Figure 9.9, P-wave velocity model perturbations obtained by Full Waveform Inversion at regional scale. The figure displays some cross-section of the 3D model. The white dot represents some of the seismic stations used in the process.

We have inverted three parameters, density, P-wave velocities and S-wave velocities. The study of these parameters is useful for understanding the geodynamics of the region. Velocity perturbations (Figure 9.9) can be interpreted in terms of geologic structures or tectonic plates. These patterns should be studied by experts in geodynamics to understand the different geological structures in the region.

**Validation.**

| Functional requirement | Target (from D5.1) | Achieved | Validated (YES/NO) |
|---|---|---|---|
| Complex physics | Using visco-elastic rheology | Multiparametric inversion for anisotropic configuration and complex rheology in use-case#1 | YES |
| Resolution. | Offshore exploration geophysics case up to 12Hz | up to 9 Hz (use-case#1). 12 Hz would likely be achieved if more computing resources had been available. | almost |
| Check-pointing | Storing iterative procedure steps | | YES |

| | | | |
|---|---|---|---|
| Data formats. | 3 components velocity dataset | 28 millions, 3-component seismograms used in use-case#1 | YES |
| Number of GPUs per run | thousands | 1024 | YES |

Table 9.3. Validation criteria for PD9.

**Involvement of the end-users**

Use-case #2 has been developed together with seismologists from GET, Midi Pyrénées Observatory (FR).

**Impact**

We tested our seismic tomography workflow on two real use cases at different scales. The more resource-intensive case is the geophysical exploration case where the large amount of data and high frequency computations require the use of PRACE Tier-0 resources. The regional scale case can be handled on Tier-1 supercomputers, but the data preprocessing is more complex in this case. It is important to make a correct selection of the data for the inversion. For this purpose, we have developed two graphical interfaces adapted to each case. The inversion workflow remains the same in both cases. We obtained new tomographic images in both cases. These models must then be interpreted by specialists (TRL=6), i.e. scientists specialized in geology or geodynamics. To this end, we propose **four output formats** for these models; the **spectral element mesh**, on a regular 3D grid, in **VTK** (or **VTU**) format or in **HDF5**. This will facilitate the use of these models for future studies.

## PD12. High-resolution volcanic ash dispersal forecast

| **PD12** | **High-resolution volcanic ash dispersal forecast** |
|---|---|
| **Leader** | Arnau Folch (BSC) |
| **Participants** | Leonardo Mingari (BSC)<br>Sara Barsotti (IMO)<br>Antonio Costa, Laura Sandri, Giovanni Macedonio (INGV) |
| **Workflow** | Ensemble-mode execution; PDAF |
| **Engine** | FALL3D (versions 8.1 and 8.2) |
| **TRL initial** | 3 |
| **TRL target** | 7 |
| **TRL achieved** | 8-9 |

### HPC Products (available software and workflows)

FALL3D is an Eulerian model for the atmospheric transport and ground deposition of volcanic tephra (ash). Versions 8.1 and 8.2 (https://gitlab.com/fall3d-distribution/v8), developed in ChEESE and associated workflows are already operative. These include generation of ensembles that can be used for probabilistic hazard assessment (PD6) and, in urgent computing mode, for ash cloud forecasting (PD12).

- Version 8.1 includes new model pre-process tasks to generate ensemble members from an unperturbed reference member, and post-process tasks to merge the single-member simulations and validate model forecasts against satellite-based and ground deposit observations. Ensemble members run concurrently in parallel, making use of a top-level hierarchy of MPI communicators involving the master ranks of each ensemble member (Folch et al., 2021).

- Version 8.2 includes an ensemble-based Data Assimilation (DA) method coupling the FALL3D dispersal model and the Parallel Data Assimilation Framework (PDAF; http://pdaf.awi.de/trac/wiki), an open-source software environment for ensemble data assimilation providing fully implemented and optimised data assimilation algorithms, in particular ensemble-based Kalman filters like LETKF and LSEIK (Mingari et al., 2021).

Ensemble members are a single realization of the model with a given set of input parameters. Ensemble-based simulations combine many ensemble members to cover and account for the whole range of uncertainties in input parameters. Two types of HPC products can be generated from PD12 ensemble-based model runs:

- Deterministic products give a deterministic forecast based on some combination of the ensemble members , e.g. on ensemble-mean tephra column mass load (g/m2), ensemble-mean tephra concentration (mg/m3) at flight levels FLs, ensemble-mean tephra deposit thickness (mm), ash cloud top-height (km a.s.l.).
- Probabilistic products give a probabilistic forecast based on counting how many of the ensemble members verify a certain condition, e.g. the exceedance of a concentration threshold, the probability of ash fallout on the ground exceeding a certain thickness, etc. By construction, ensemble members are generated by sampling a user-defined Probability Density Function for each input variable, i.e. assigning more members to more likely values. As a result, all members have an equal weight in the final merging stage.

**Use case #1. The VOLCICE-2021 exercise**

On 12 March 2021, a VOLCICE exercise was scheduled to practice the response to an eventual explosive eruption at Beerenberg volcano (Jan Mayen, Norway). The exercise is part of the VOLCICE series played by the Icelandic Meteorological Office (IMO) in collaboration with London Volcanic Ash Advisory Center (VAAC) and ISAVIA (the air navigation service provider in Iceland). The Jan Mayen exercise also engaged the Norwegian Meteorological Institute (MET Norway).

***Summary of technological achievements (HPC performance, etc)***

Table 12.1 compares the current operational products (London VAAC) with the ChEESE PD12.

| | London VAAC | ChEESE PD12 |
|---|---|---|
| **Ash dispersal model** | NAME, 1 single run | FALL3D (v8.1), 21 ensemble members (*) |
| **Meteorological model** | U.K. UM (10 km resolution) | GFS + WRF ( 5km resolution) |
| **Output grid resolution** | 40 km, 3 thick layers | 5 km, 10 levels (Flight Levels) |
| **Output time resolution** | 6 h (T+0, T+6, T+12, T+18) | 1 h (from T+0 to T+24 every 1h) |
| **Deterministic forecast products** | VAG (qualitative). No-fly zone at 3 thick layers | Concentration at 10 FLs (with 0.2, 2 and 4 mg/m3 contours)<br>Ash cloud Column mass (g/m2)<br>Ash cloud top (km a.s.l.)<br>Ground deposit thickness (mm) |
| **Probabilistic forecast products** | No | Prob. of no-fly zone areas at 10 FLs<br>Prob. of cloud (column mass) detection<br>Prob. of ash fallout at ground |

Table 12.1. Summary of scientific/technological advances of ChEESE PD12 with respect to state-of-the-art methodologies. (*) using 84 nodes (1344 cores) of the nord3

machine (former MN-3) at BSC. Each ensemble member runs on 4 nodes (64 cores). In addition, backup (redundant) runs were ready on Irene-rome at TGCC.

***Scientific achievements***

When compared with the London VAAC, the PD12 forecasts have (Figure 12.1):

- Slightly better time latency (21 vs 26 min)
- Much higher space-time resolution (5 vs 40 km and 1 vs 6h)
- Array of quantitative products (including probabilistic forecasts)

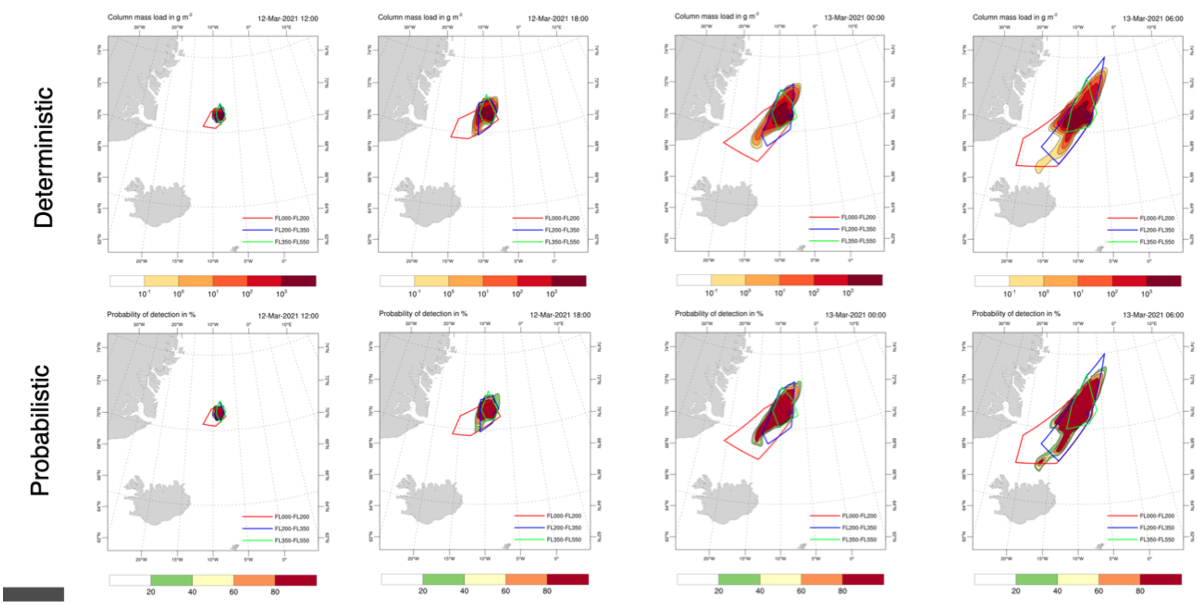


Figure 12.1. Comparison between deterministic VAAC results (contour lines at 3 different FL intervals) and the PD12 deterministic (top row) and probabilistic forecasts (bottom row) for the VOLCICE exercise.

**Use case #2. La Palma eruption**

During the eruption from Cumbre Vieja volcano at La Palma island (starting on 19th September 2021, still on-going at the time of submitting this report), the scientific committee of the PEVOLCA (in Spanish "Plan Especial de Protección Civil y Atención de Emergencias por riesgo volcánico en la Comunidad Autónoma de Canarias") asked for an operational service for ash/SO2 forecast. From 25th September 2021, PD12 is run daily at the MN-4 supercomputer and the resulting forecast products are used by the PEVOLCA committee to advise authorities on regional airport management and air traffic. Every morning the PEVOLCA delivers results to ENAIRE, the national air traffic regulator that decides on airport/airspace closure. Even with limited computational resources, this use case is testing PD12 at a TRL of 9.

***Summary of technological achievements (HPC performance, etc)***

Two ensemble-based forecasts run daily on two computational domains, one for the Canary Islands and one regional, including NW Africa and the Iberian peninsula.

| | Canary Islands domain | Regional domain |
|---|---|---|
| Ash dispersal model | FALL3D (v8.1), 6 ensemble scenarios (*) | FALL3D (v8.1), 6 ensemble scenarios (*) |
| Output grid resolution | 2 km, 40 vertical levels | 10 km, 40 vertical levels |
| Output time resolution | 1 h (from T+0 to T+36 every 1h) | 1 h (from T+0 to T+36 every 1h) |
| Forecast products | Concentration at 10 FLs (with 0.2, 2 and 4 mg/m3 contours)<br>Ash cloud Column mass (g/m2)<br>Ash cloud top (km a.s.l.)<br>Ground deposit thickness (mm) | Concentration at 10 FLs (with 0.2, 2 and 4 mg/m3 contours)<br>Ash cloud Column mass (g/m2)<br>Ash cloud top (km a.s.l.)<br>Ground deposit thickness (mm) |

Table 12.2. Simulation parameters. (*) 1536 cores of MN-4 machine at BSC, time latency 15 min.

***Scientific achievements***

Service already operational (TRL 9), used for decision making by ENAIRE and airlines (see Figure 12.2 and the press note at https://geo3bcn.csic.es/index.php/news-events/news/1916-geo3bcn-y-el-bsc-trabajan-en-la-prediccion-de-las-nubes-de-ceniza-y-so2-del-volcan-de-la-palma?jjj=1634457146015 )

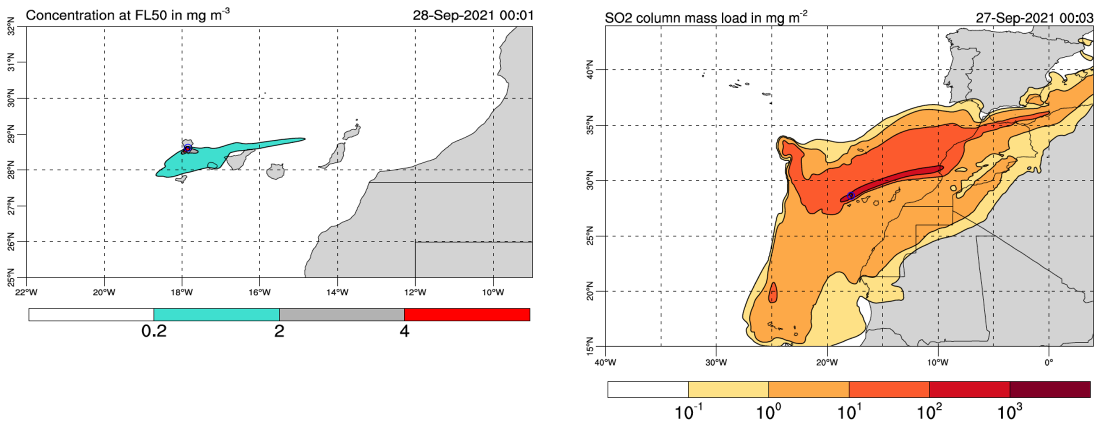


Figure 12.2. Example of a PD12 deterministic forecast at La Palma Island for the local (left) and regional (right) domains. Results correspond to 26th Sep. 2021 and show concentration at FL050 and SO2 column mass load.

***Scientific Products (including publications)***

Reports are confidential

**Use case #3. ICAO-FICTITUS exercise**

On 9th December 2021, the International Civil Aviation Organization (ICAO) will organise a FICTITUS exercise in South America, under the area of responsibility of the Buenos Aires VAAC. In this exercise ensemble forecasts of volcanic ash will be generated using the FALL3D model considering a hypothetical eruptive event within the area of responsibility of the Buenos Aires Volcanic Ash Advisory Center (VAAC). Numerical simulations will be performed simultaneously in the Marenostrum 4 at the Barcelona Supercomputing Center (BSC) and at the National Meteorological Service (Argentina) by the Buenos Aires VAAC. The objective of this exercise is to practice the response to reports of volcanic ash within the region of responsibility of the Buenos Aires VAAC in an operational environment and provide ensemble forecasts to the aviation community of ash cloud extent and movement. The ash warnings describe the current and future extent of volcanic ash clouds and can be issued in the form of Volcanic Ash Advisories (VAA) and SIGMETs messages. Results will be included in the final report.

**Validation.**

| **Functional requirement**<br><br>**(these are examples)** | **Target (from D5.1)** | **Achieved** | **Validated (YES/NO)** |
|---|---|---|---|
| Time to solution | 1 hour maximum | 15 to 26 min | YES |
| Resolution. | 10 km (regional domains) to 5 km (local domains) | 10 km (regional domains), 2 km (local domains) | YES |
| Number of simulations. | 1 per scenario | Depends on resources | YES |
| Data formats. | netCDF, maps | netCDF, maps | YES |
| Uncertainty quantification. | Give measure of forecast uncertainty | Yes, based on ensemble spread | NO |

Table 12.3. Validation criteria for PD12.

**Involvement of end-users**

Good engagement with IUB members and other end-users in the validation (and use) of PD12, including:

- Air traffic regulators: ISAVIA in Iceland and ENAIRE in Spain
- Volcanic Ash Ash Advisory Centers (Buenos Aires)
- Meteorological Services (Iceland, Norway, Spain)

- Civil protection agencies
- Other IUBs: IGN, Mitiga

**Impact**

High. Current operational products are at coarse space-time resolution. PD12 effectively tested at 2 to 10 km resolution with run forecasts within 1 hour whenever substantial changes occur in the eruption conditions. Forecast window up to 48 hours to cover tactical and pre-tactical flight design phases. In addition, decision makers have a forecasted uncertainty quantification indicator and probabilistic products obtained from perturbations of ensemble members and/or scenarios.

***Scientific Products***

- Folch, A., Mingari, L., Prata, A., Ensemble-based forecast of volcanic clouds using FALL3D-8.1, Frontiers, under review, 2021.
- Mingari, L., Folch, A., Prata, A. T., Pardini, F., Macedonio, G., and Costa, A.: Data Assimilation of Volcanic Aerosols using FALL3D+PDAF, Atmos. Chem. Phys. Discuss., https://doi.org/10.5194/acp-2021-747, in review, 2021.